

<p>Volume 75, Issue 2, August 2010</p> <p>Founding Editors: RICHARD H. DAY SIDNEY G. WINTER</p> <p>Honorary Editors: GEORGE A. ANGEL OF RICHARD A. EASTERLUN ALAN KIRMAN REINHARD SELTEN VERNON L. SMITH OLIVER E. WILLIAMSON</p> <p>Editor: J. BARKLEY ROSSER, JR.</p> <p>Co-Editors: CATHERINE ECHEL GEORG KIRCHSTEIGER CLAUDE MENARD JOHN OCHSNER NICOLAAS J. VRIEND</p> <p>Managing Editor: DEBRA DOVE</p> <p>Book Review Editor: ROSEF KOYPL</p> <p>Associate Editors: JOHN A. CANTWELL JEFFREY P. CARPENTER GARY CHARNESS SHUANG CHEN ANDREW CLARK JOHN COLUCCI RACHEL COXSON SATHYA P. DAS HERBERT DAMID DAVID DEQUECH ZENGRUO M. DUFWENBERG D. ENGELBACH ENRIQUE FADAS DAMIEN FILSON DUNCAN K. FOLEY SIMON GAECHTER LAURA GASTONI HERBERT GINTIS CAROL GRAHAM WILLIAM C. GRANT DAVID GREYHER GEOFFREY M. HODGSON CARB HOMMES ED HOPKINS STEVEN HORWITZ DANIEL HOUBER STEFFEN HUCK PENKA IMAMKURBAS GIULIA IORI MICHAEL KOPPEL TIBUR KUJRAJ CHRISTOPH KUZMICKS JOSH LERNER GARY LIRICAP JOHN A. LIET THOMAS LUX INES MACHOINZLER ANDO MATSUOCTO WILLIAM NELSON CHARLES NOUSGAR KARINE NYBORG AXEL OCKENFELS MARK PHILLE VICTOR RAJ RAO JASON POTTS LUCA RICCIOTTI ABDOLKARIM SAGHIAN RAJIV SETHI JASON F. SHOUBEN PETER SWOTT STEPHEN C. SMITH RINE STENBACKA MARE THURSBY MARIE CLAIRE WILLEVAL BAUKE VISSER FRANK WESTERHOFF WILLIAM C. WOOD</p>	<p>ISSN 0167-2681</p> <h1 style="text-align: center;">JOURNAL OF Economic Behavior & Organization</h1> <table border="0" style="width: 100%;"> <tr> <td>D.K. Foley, What's wrong with the fundamental existence and welfare theorems?</td> <td style="text-align: right;">115</td> </tr> <tr> <td>M. Dufwenberg, R. Sunderam and D.J. Butler, Epiphany in the Game of 21</td> <td style="text-align: right;">132</td> </tr> <tr> <td>U. Gneezy, A. Rustichini and A. Vostroknutov, Experience and insight in the Race game</td> <td style="text-align: right;">144</td> </tr> <tr> <td>D. Di Cagno and E. Scubba, Trust, trustworthiness and social networks: Playing a trust game when networks are formed in the lab</td> <td style="text-align: right;">156</td> </tr> <tr> <td>J.W. Boudreau, Stratification and growth in agent-based matching markets</td> <td style="text-align: right;">168</td> </tr> <tr> <td>O. Gürtler and M. Kräkel, Optimal tournament contracts for heterogeneous workers</td> <td style="text-align: right;">180</td> </tr> <tr> <td>T. Feldman, Portfolio manager behavior and global financial crises</td> <td style="text-align: right;">192</td> </tr> <tr> <td>H. Dawid, M. Kopel and P.M. Kort, Innovation threats and strategic responses in oligopoly markets</td> <td style="text-align: right;">203</td> </tr> <tr> <td>C.M. Capra, K.F. Lanier and S. Meer, The effects of induced mood on bidding in random nth-price auctions</td> <td style="text-align: right;">223</td> </tr> <tr> <td>P. Courty and M. Pagliero, Price variation antagonism and firm pricing policies</td> <td style="text-align: right;">235</td> </tr> <tr> <td>H. Ku and A. Zussman, Lingua franca: The role of English in international trade</td> <td style="text-align: right;">250</td> </tr> </table> <p style="text-align: center;"><i>(Contents continued on outside back cover)</i></p> <p style="text-align: center;">www.elsevier.com/locate/jebo</p>	D.K. Foley, What's wrong with the fundamental existence and welfare theorems?	115	M. Dufwenberg, R. Sunderam and D.J. Butler, Epiphany in the Game of 21	132	U. Gneezy, A. Rustichini and A. Vostroknutov, Experience and insight in the Race game	144	D. Di Cagno and E. Scubba, Trust, trustworthiness and social networks: Playing a trust game when networks are formed in the lab	156	J.W. Boudreau, Stratification and growth in agent-based matching markets	168	O. Gürtler and M. Kräkel, Optimal tournament contracts for heterogeneous workers	180	T. Feldman, Portfolio manager behavior and global financial crises	192	H. Dawid, M. Kopel and P.M. Kort, Innovation threats and strategic responses in oligopoly markets	203	C.M. Capra, K.F. Lanier and S. Meer, The effects of induced mood on bidding in random n th-price auctions	223	P. Courty and M. Pagliero, Price variation antagonism and firm pricing policies	235	H. Ku and A. Zussman, Lingua franca: The role of English in international trade	250
D.K. Foley, What's wrong with the fundamental existence and welfare theorems?	115																						
M. Dufwenberg, R. Sunderam and D.J. Butler, Epiphany in the Game of 21	132																						
U. Gneezy, A. Rustichini and A. Vostroknutov, Experience and insight in the Race game	144																						
D. Di Cagno and E. Scubba, Trust, trustworthiness and social networks: Playing a trust game when networks are formed in the lab	156																						
J.W. Boudreau, Stratification and growth in agent-based matching markets	168																						
O. Gürtler and M. Kräkel, Optimal tournament contracts for heterogeneous workers	180																						
T. Feldman, Portfolio manager behavior and global financial crises	192																						
H. Dawid, M. Kopel and P.M. Kort, Innovation threats and strategic responses in oligopoly markets	203																						
C.M. Capra, K.F. Lanier and S. Meer, The effects of induced mood on bidding in random n th-price auctions	223																						
P. Courty and M. Pagliero, Price variation antagonism and firm pricing policies	235																						
H. Ku and A. Zussman, Lingua franca: The role of English in international trade	250																						

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Epiphany in the Game of 21

Martin Dufwenberg^{a,b,*}, Ramya Sundaram^c, David J. Butler^d^a Department of Economics, University of Arizona, Tucson, AZ 85750, USA^b Department of Economics, University of Gothenburg, Sweden^c Department of Economics, George Washington University, 2115 G Street, NW, Washington, DC 20052, USA^d Economics Program, UWA Business School, University of Western Australia, 35 Stirling Highway, WA 6020 Australia

ARTICLE INFO

Article history:

Received 17 May 2009

Received in revised form 26 March 2010

Accepted 29 March 2010

Available online 7 April 2010

JEL classification:

C7

C9

Keywords:

Game theory

Reasoning

Epiphany

Mimicking

Dominant strategies

Learning

ABSTRACT

Does performance in strategic settings depend on whether players realize that an optimal way to play is feasible? We introduce a zero-sum game of perfect information, simple enough to allow computation of optimal play yet sufficiently complicated that most participants initially fail. This borderline solvability-by-humans makes it a suitable research tool for experimentally evaluating if play is affected by whether it dawns on a subject that an analytic solution may be possible. Our design includes a way to control for such insight. We also examine how learning transfer across games affects subsequent learning towards optimization. Applications include the facilitation of learning how to plan ahead when actions are needed today but the consequences are temporally distant.

© 2010 Elsevier B.V. All rights reserved.

Epiphany –

a sudden manifestation or perception of the essential nature or meaning of something [Webster's Dictionary]

1. Introduction

Understanding the workings of the human mind can be crucial to economists. Economic outcomes depend on behavior, and behavior is shaped by how people reason. In order to make reliable predictions it is useful to know the type of reasoning triggered by various situations and the effect on behavior and outcomes.

Our paper contributes to the literature which seeks related insights through experimental games.¹ We focus on an old but mostly little-known game that we call *The Game of 21*; G_{21} for short. A variant of this game was used in the ‘immunity challenge’ of an episode of the TV series ‘Survivor’, as discussed in Dixit (2005). The rules: Two players, call them *White* and *Green*, take turns. *White* begins. To start off, he can choose either 1 or 2. *Green* observes this choice, then increments the “count” by adding one or two. That is, if *White* chooses 1 *Green* can follow up with 2 or 3; if *White* chooses 2 *Green* can follow up with 3 or 4. *White* then observes *Green*’s choice, and again increments the count by adding one or two. The

* Corresponding author at: Department of Economics, University of Arizona, Tucson, AZ 85750, USA. Tel.: +1 520 626 1540; fax: +1 520 621 8450.

E-mail addresses: martind@eller.arizona.edu (M. Dufwenberg), ramya.sundaram@gmail.com (R. Sundaram), david.butler@uwa.edu.au (D.J. Butler).

¹ For an entry, see the paper by Nagel (1995) which introduced so-called guessing games, the survey in chapter 6 in Camerer’s (2003) book *Behavioral Game Theory* which covers many other games, and our further discussion and references in Section 2.

game continues with the players taking turns, each player incrementing the count by one or two. The player who reaches 21 wins.

We now invite our readers to ask themselves, before reading on, what they would do in this game.

G_{21} features a second-mover advantage. Any strategy which whenever possible selects a multiple-of-three is dominant for *Green*.² Such a strategy ensures that along the path of play *Green* makes the choices 3, 6, . . . , 18, 21, and thus guarantees victory. Furthermore, in any sub-game where a co-player has not chosen a multiple-of-three in the preceding stage, the player to move may use a dominant strategy for that sub-game involving choices of multiples-of-three from that point on.

Did you figure this out? Our experience says many people don't grasp it, including "professional" conference audiences. Dixit (2005) recounts the failure of participants in the reality-TV show 'Survivor' in figuring out the dominant strategy in a variant of G_{21} . In the '21 Flags Game,' aired on episode 6 of 'Survivor Thailand,' each of two teams could remove 1, 2 or 3 flags each turn in their quest to remove the last flag. The team that lost this game had to vote off one of its members, decreasing its chance of winning the eventual million-dollar prize. Dixit notes that players got almost every choice wrong; they should have aimed to leave the other tribe with a multiple-of-four.

Why this difficulty in figuring out a finite two-player zero-sum game with perfect information? On further reflection one realizes there may be (at least) two reasons: *First*, a player may not realize the analytical nature of the problem; our experience using G_{21} for teaching suggests even economics students often fail to realize that an analytic solution is possible. Once the solution to this game is known, it seems difficult to fathom that people may not realize an analytic solution is possible. But bear in mind that most situations in life lack dominant strategies. Depending on a subject's associations, he or she may not think of the possibility that an optimal way to play G_{21} could exist.³ *Second*, even if one realizes that logical analysis may hold the key, finding the answer may prove too difficult as it requires going through several steps-of-reasoning.⁴

Epiphany! That's what a player needs to master G_{21} . To study how players may achieve epiphany, we introduce a second, but shorter game: *The Game of 6*, G_6 for short, is played the same way as G_{21} except whoever reaches 6 wins. Try G_6 on anyone, and they quickly figure out that they can win by picking 3 as *Green*.

How might playing G_6 facilitate learning to play optimally in G_{21} ? There are two plausible learning mechanisms—(a) mimicking: i.e., a player learns that playing a multiple of 3 is a winning strategy in G_6 , and may simply imitate this strategy in the longer G_{21} game, or (b) playing G_6 makes players realize that an analytic solution is possible in G_{21} . We could call this a mini-epiphany. Next, s/he discovers and understands the dominant strategy. We call this full epiphany, or complete epiphany.⁵

Whatever the form of learning across games, if any learning occurs, then a person playing G_6 is likely to carry some insight of an analytical solution over to G_{21} . Cooper and Kagel (2008, 2009) offer supporting evidence for the possibility of learning transfer across games when the context is similar, and G_6 and G_{21} are similar except that the latter game involves much longer paths of play. They also note that Fudenberg and Kreps (1988) suggest ". . . players infer about how their opponents will act in one situation from how opponents acted in other, similar situations".

We have two treatments. In the first subjects first play G_{21} several times and then G_6 several times, in the other the order is reversed (G_6 before G_{21}). We ask: will subjects playing G_{21} in the latter treatment play better or learn faster than subjects playing G_{21} in the former treatment? This is our first research question.

Our second research question concerns whether and how subjects arrive at complete epiphany, i.e., how do they discover the dominant strategy in G_{21} , given that they have already played G_6 . Is it the case that, over time and as subjects play more and more games, they learn gradually in the sense that they choose multiples-of-three at incrementally lower counts in G_{21} (epiphany by the backdoor)? Or could it be that subjects show no evidence of gradual learning before epiphany occurs (learning with a leap)? We study the patterns, focusing on the data from the treatment where subjects play G_6 before G_{21} .

How does this approach add to the previous literature on strategic reasoning in games? We answer this question in Section 2, as we review related literature. Thereafter, Section 3 describes our experimental design, Section 4 reports results regarding our two research questions, and Section 5 concludes.

² *Green* has many dominant strategies since he may make any choice after histories where he has previously not responded with multiple-of-three at some point (such histories will, of course, be counter-factual if *Green* follows a dominant strategy). The general insight that some player in G_{21} must have a dominant strategy can be gleaned (on a little reflection) by abstract principles (that G_{21} is a finite two-player zero-sum two-outcome games with perfect information) from Ewerhart (2000).

³ Or consider G_{21} with three rather than two players taking turns; no dominant strategy exists in this modified game. Is it really so obvious when to look for, or not conceive of, dominant strategies?

⁴ First, and trivially, realize that a choice of 21 wins. Second, realize that if one chooses 18 then a win can be guaranteed. Third, realize that if one chooses 15 then one can similarly secure 18, and so on. Ultimately, if one chooses 3, and then a multiple-of-three in every subsequent move, then one can secure a win. According to the implicitly suggested metric, this calculation requires six steps of reasoning. Note that the described process resembles backward induction, but in fact is not backward induction since no reference is made to optimal subsequent co-player choices. The process considers each player i in isolation and works backwards on i 's nodes assigning an optimal choice only if this can be done *regardless* of subsequent opponent choices, and so exhibits non-existence except if a dominant strategy is uncovered for each sub-game.

⁵ Our initial approach to this research project had been to study the second approach to learning: i.e., to induce mini-epiphany in subjects (using G_6) and then to understand whether this helped trigger complete epiphany, and so on. Both our referees and co-editor brought up the issue that, while that was our intended purpose, the results we observed could also be explained by the hypothesis that players were mimicking the multiple-of-3 strategy that they learnt in the shorter game when playing the longer game. We think this is a fair argument, and have included, ex-post, the learning mechanism of mimicking. However we present evidence at the end of Section 4.1 that the type of learning we observe seems to be through the 'mini-epiphany and complete epiphany' mechanism, rather than the mimicking mechanism.

2. Related literature

To see how we add to preceding literature, let us first describe a version of the classical guessing game: $N > 2$ players simultaneously pick numbers in the range $[0,100]$. Whoever is closest to $2/3$ of the average wins/splits a prize. The unique Nash equilibrium (also the result of iterated elimination of weakly dominated strategies) is for each player to pick 0. However, in experiments, choices are all over, 0s are rare, and 0s never win (unlike choices around 20); see e.g. Nagel (1995) and Camerer (2003, chapter 5).

This is sometimes taken to illustrate subjects' bounded reasoning abilities. High choices certainly make it clear that the players *collectively* do not manifest the degrees of mutual beliefs about mutual beliefs. . . about rational choices that might correspond to various rounds of iterated dominance. This does not, however, reveal much about any *individual's* ability to reason deeply. A smart and potentially deep-reasoning individual should avoid the equilibrium strategy of 0 since most of the others choose high numbers!⁶

The game we study avoids such interpretational ambiguities. Playing the dominant multiples-of-three strategy is a best response regardless of beliefs about others. Failure to choose a feasible multiple-of-three very likely indicates failure to work out the dominant strategy. Moreover, we can infer something regarding the number of steps-of-reasoning a subject is capable of by observing how early in the count of a game he starts choosing multiples-of-three (cf. footnote 4).

Another contribution of ours is best understood with reference to recent work on cognitive hierarchy or level- k models which can account for subjects' play in many experiments using simultaneous-move games. The key idea, pioneered by Nagel (1995) and Stahl and Wilson (1994, 1995),⁷ is that players are heterogeneous in terms of strategic sophistication. For example, level-0 players may choose randomly across all strategies. Level-1 players assume everyone else is level-0, and best respond; level-2 players assume everyone else is a level-1 player, and best respond; etc.⁸ Estimations of such models, for specific games, indicate a distribution of players concentrated around small but non-zero k 's. Costa-Gomes and Crawford (2006) report that "many subjects' systematic deviations from equilibrium can be confidently attributed to non-equilibrium beliefs rather than irrationality" (p. 1767), thus describing data from games where subjects presumably succeeded in optimizing given their beliefs of others' strategies.

But this is in contrast with recent findings by Grosskopf and Nagel (2008), on guessing games with $N=2$. Two players simultaneously pick numbers in the range of $[0,100]$ and whoever is closest to $2/3$ of the average number wins/splits a prize. The change from $N > 2$ to $N=2$ alters the game's properties: a choice of 0 is now dominant. Student subjects as well as professional audiences at economics and psychology conferences made choices that were not significantly different from the choices made in $N > 2$ treatments. With $N=2$, 90% of the students and 63% of the professionals chose a dominated strategy! If one were to apply a level- k model with a distribution of players concentrated around small but non-zero k 's it would suggest most players (all those for whom $k > 0$) should choose 0, at odds with Grosskopf and Nagel's data.⁹

Let's take stock. It seems that in some games it is easier for subjects to optimize than in others. We face the challenge of explaining how subjects calculate and learn what is in their best interest. This is a largely open research area, and we take early steps of exploration.¹⁰ G_{21} joins Grosskopf and Nagel's $N=2$ games in having a dominant strategy which is non-obvious to compute. Our approach is different from theirs (as well as from almost all studies of level- k play) in that we have a game with a sequential structure rather than simultaneous moves. This feature enables us to get some insights regarding how close subjects come to optimizing (cf. footnote 4). Another difference is that we focus on learning and how learning is transferred across games, whereas most of the other work we cited looks at games played only once or games where subjects get no feedback between rounds.

After we started our project we learnt of work by Gneezy et al. (this issue), involving similar games, conducted independently. G_{21} features counting to 21 in increments of one or two; Gneezy et al. have players count to 15 (or 17) with steps of one to three (or four). Some patterns of play accord well across studies, but research questions differ. Gneezy et al. do not consider our key notion of inducing epiphany through playing a much simpler game before playing a more challenging game. Rather they focus on patterns of learning across rounds of play of a given game, or games which are similar in terms of path-length. On the other hand, Gneezy et al. explore issues concerning response times and talk-aloud protocols which we do not consider.

⁶ Camerer (2003, p.17) recognizes this confound and recounts how one player he knew to be very clever chose 18.1. Asking him later to explain his choice, he said he knew 0 was the equilibrium but believed his colleagues (all were Board members at Caltech) would only average two steps of reasoning and pick 25. He optimized on that assumption, adding a little extra in case an odd high number were also chosen.

⁷ For further developments or applications, see Bosch-Domenech et al. (2002), Camerer et al. (2004), Costa-Gomes and Crawford (2006), Costa-Gomes et al. (2001), Crawford (2003), Crawford et al. (2008), Crawford and Iriberry (2007a, 2007b), Gneezy (2005), Ho et al. (1998), Östling et al. (2009), and Selten et al. (2003).

⁸ Some versions allow that level- k players best respond to some combination of players at level- k' , for $k' = 0, 1, \dots, k-1$. See e.g. Camerer et al.

⁹ In follow-up research, Chou et al. (2009) suggest that the failure of subjects to choose optimally in the 2-person guessing game may rest with the failure of the experimenter to design instructions and procedures that allow subjects to properly grasp the game. They go to great lengths to achieve such control, and nevertheless meet with only limited success.

¹⁰ A different approach is developed by Johnson et al. (2002) who employ the 'Mouselab' system to study patterns of information search in alternating-offer shrinking pie games. They report that players tend not to backward induct; indeed a minority did not even glance at the pie sizes in later rounds, making backward induction impossible. Gabaix and Laibson (2000) develop a forward-looking algorithm for solving complex decision trees without invoking backward induction.

McKinney and Van Huyck (2006, 2007) study depth-of-strategic-reasoning related issues in Nim, an ancient game named in modern times by Bouton (1901–02). Again some features of play accord between studies, but G_{21} and Nim are sufficiently different that a direct comparison is difficult.¹¹ McKinney and Van Huyck also put more emphasis on identifying bounds of human reasoning; again they do not deal with whether playing a simpler game may induce faster learning in a more challenging version.

3. Design

Our subject pool was unusual. One of us was teaching two sections of intermediate microeconomics. The course involved discussion of experimental methodology and results. To get the students excited about the topic, they were promised in-class experience of a “real” experiment, one generating data meant for publication. After some negotiation the Human Subjects Protection Program of the University of Arizona gave permission.¹² Sessions were conducted at the Economic Science Laboratory. Since subjects were in class, we had no reason to make sure each was compensated for their time. We used a pay-a-random-subset-of-subjects approach as advocated by Bolle (1990): two subjects from each treatment were selected at random (one for G_{21} ; one for G_6) and paid \$5 for each game won.

We had two treatments: in the G_{21} -then- G_6 treatment subjects first played G_{21} five times and then G_6 five times, in the G_6 -then- G_{21} treatment the order was reversed. Subjects were not permitted to communicate with each other once the experiment had commenced, other than through selecting their choices of integers.

The G_{21} -then- G_6 treatment had 42 participants comprising seven groups of six subjects. Each subject received a “player ID” (A, B, C, D, E, or F), read through a “subject disclaimer form”, and then got instructions with the rules of G_{21} . Each pair of members of each group played G_{21} once, with new matches proceeding round-robin style with players alternating between *White* and *Green* positions.¹³ Play began once subjects had spent sufficient time studying the rules of the game. No hard time limit was imposed. Game sheets for each round were collected only after the last pair of players in that round had finished playing.

We thus had seven groups, each featuring five rounds of play of G_{21} , with three games (each with two players) per round. After all of these $7 \times 5 \times 3 = 105$ games had terminated,¹⁴ with a winner determined for each, instructions describing the rules of G_6 were distributed, and round-robin play ensued as before, with another $7 \times 5 \times 3 = 105$ games. The G_6 -then- G_{21} treatment had the same format, except the order of the games was reversed. We had 30 participants, producing 5 groups. We thus had $5 \times 5 \times 3 = 75$ games of G_6 and another $5 \times 5 \times 3 = 75$ games of G_{21} in this treatment.

4. Results

Epiphany in G_{21} involves the dawning on a player that s/he may have a way of playing that guarantees a win. This is a cognitive concept which we can only study indirectly. We use G_6 as a tool to induce such insight in our subjects. The idea is that once subjects figure out that G_6 is “solvable”—either by understanding the analytic nature of the game (i.e., mini-epiphany) or by learning that perhaps playing multiples-of-3 may work in the longer game (mimicking)—they will start thinking that G_{21} may be solvable too since the games have a similar structure.¹⁵ Hence, when we analyse our data we will assume that subjects in the G_6 -then- G_{21} treatment reach some kind of insight before playing G_{21} , while subjects in the G_{21} -then- G_6 treatment may or may not have reached any such insight before playing G_{21} . Conditional on that maintained assumption, we then test our two main hypotheses mentioned in Section 1.

¹¹ Several people suggested to us that G_{21} is a version of Nim. However, no Nim game exists which has the same extensive game form as G_{21} . While G_{21} and Nim both are finite two-player zero-sum two-outcome games with perfect information, in G_{21} the root of any sub-game (other than at the count of 20) has a binary choice set, a feature which cannot be preserved throughout any Nim game rich enough to allow as long paths of play as G_{21} requires. Moreover, Bouton’s (ingenious!) solution method, while similar to the pick-multiples-of-three solution of G_{21} in the sense that it too produces a method by which a winning position (“safe combination”) can be maintained through play, is very different in its details (which involve manipulations of binary scale of notation representations of positions) and does not apply to G_{21} . Finally, the number of moves needed to win using a dominant strategy is not as unambiguously defined in Nim, making the link to steps-of-reasoning less straightforward.

¹² The issue was that research experiments usually occur outside of class, since participation is supposedly voluntary in a way which classes are not. We got around this by providing an alternative lecture (on theory regarding the involved games) for students who wished to opt out (no one did). We acknowledge that, as a referee put it, one may “worry a bit about loss of control as students are engaged in reasonably long-run relationships”. On this dimension, the study of Gneezy et al. complements ours in that they use a more standard set of subjects and nevertheless overall patterns of play in G_{21} shares the feature (as we will see) that subjects often fail to implement a dominant strategy.

¹³ Appendix A contains instructions, game sheets, and the schedule-cards/protocol for matching pairs of players across rounds (which followed a so-called *Howell movement*, commonly used for conducting contract bridge-pairs tournaments). The format of the game sheets we used was chosen to visually facilitate understanding the nature of the game and thus epiphany, at least compared with methods where the sequence of choices must be remembered mentally.

¹⁴ Our data analysis, however, is based on only 104 of these games. One pair of subjects (round 1, group 4, players E and F) had not understood the instructions and played erroneously in their first round.

¹⁵ This need not mean that they figure out what the dominant strategy in G_{21} is, only that they will realize that it may make sense to look for a dominant strategy in G_{21} .

Table 1
Perfect Play in G_6 .

Round	G_6 -then- G_{21}		G_{21} -then- G_6		Pooled	
	Perfect games	Percentage	Perfect games	Percentage	Perfect games	Percentage
1	12/15	80	18/21	86	30/36	83
2	14/15	93	20/21	95	34/36	94
3	13/15	87	21/21	100	34/36	94
4	15/15	100	20/21	95	35/36	97
5	14/15	93	20/21	95	34/36	94
All rounds	68/75	91	99/105	94	167/180	93

Notes: Columns 2, 4, and 6 (“Perfect Games”) list the relative number of G_6 where Green played perfectly, for each round, by treatment and pooled. Columns 3, 5, and 7 provide the associated percentages.

Table 2
Perfect play in G_{21} .

Round	G_6 -then- G_{21}		G_{21} -then- G_6		Pooled	
	Perfect games	Percentage	Perfect games	Percentage	Perfect games	Percentage
1	3/15	20	2/20	10	5/35	14
2	5/15	33	3/21	14	8/36	22
3	6/15	40	4/21	19	10/36	28
4	5/15	33	7/21	33	12/36	33
5	8/15	53	6/21	29	14/36	39
All rounds	27/75	37	22/104	21	49/179	27

Notes: Columns 2, 4, and 6 (“Perfect Games”) list the relative number of G_{21} where Green played perfectly, for each round, by treatment and pooled. Columns 3, 5, and 7 provide the associated percentages.

This approach is meaningful only if two *preliminary results* hold:

PR(i) Most subjects playing five rounds of G_6 realize that G_6 may be solvable by rational calculation. (If they did not, our idea that such an insight extends to G_{21} would lose its basis.)

PR(ii) Most subjects playing the *Green* position in G_{21} for the first time do not immediately figure out that choosing the multiples-of-three is the best they can do. (If they did, then our conjecture that subjects in the G_{21} -then- G_6 treatment may not have reached complete epiphany before playing G_{21} would be vacuous.)

This section has three subsections: in subsection (4.1) we establish preliminary results PR(i) and PR(ii), and in Sections 4.2 and 4.3 we consider our two main research hypotheses.

4.1. Two preliminary results

PR(i) above is supported: most subjects playing five rounds of G_6 realize that G_6 may be solvable by rational calculation. Table 1 shows this with data from G_6 for both treatments. 167 of 180 *Green* players (93 percent) play perfect games—that is, their first move is 3, and their second move is 6, at which point they win.¹⁶

PR(ii) is supported too: most subjects playing the *Green* position in G_{21} for the first time do not immediately figure out that choosing multiples-of-three is the best they can do. The evidence is in Table 2. Across treatments, in G_{21} , only 49 of 179 games (27%) are played perfectly.¹⁷ The rates of perfect play are especially low in the early rounds of the G_{21} -then- G_6 treatment (e.g. 2 out of 20, or 10%, in round 1).

One final comment about PR(i) and PR(ii): In G_6 a player might stumble on his optimal strategy serendipitously – if he flipped a coin he would choose 3 with probability 1/2 – and from there win almost for sure (169 out of 180 *Green* players chose 3; only two of those 169 *Green* players then failed to win in G_6). It is much harder to stumble into the optimal strategy in G_{21} . We can, however, control for this potential confound if we simply count the number of *Green* players in G_{21} who chose 3; 113 out of 179 *Green* players did so. While this proportion is significantly greater than expected from a coin-flip, it is also significantly lower than the proportion of players who chose 3 in G_6 ($z = 7.1$).¹⁸ The conclusion: more *Green* players in G_6 than in G_{21} chose 3 because they figured out their dominant strategy before choosing. Furthermore, of the 113 *Green* players mentioned above, 52 (out of 75) were from the G_6 -then- G_{21} treatment, while only 61 (out of 104) were from the

¹⁶ There is no significant difference in perfect play of G_6 between the G_6 -then- G_{21} and G_{21} -then- G_6 treatments.

¹⁷ There is a significant difference in perfect play of G_{21} between the G_6 -then- G_{21} and G_{21} -then- G_6 treatments. This finding is central to our first main hypothesis, discussed further in the next subsection.

¹⁸ Also, the difference in the proportion of persons playing perfectly in *all* rounds of G_6 (93 percent) vs the proportion playing perfectly in *all* rounds of G_{21} (27 percent) is overwhelmingly significant ($Z = 12.65$).

Table 3Comparing play in G_{21} across the two treatments.

Measure	G_6 -then- G_{21}		G_{21} -then- G_6	
	Raw count	Percentage	Raw count	Percentage
Perfect play	27/75	37	22/104	21
Median indicated rationality	6		12	
Moment of epiphany in				
Round 1	5/30	17	3/42	7
Round 2	4/30	13	3/42	7
Round 3	2/30	7	2/42	5
Round 4	2/30	7	3/42	7
Round 5	1/30	3	4/42	10
No moment of epiphany	16/30	53	27/42	64

Notes: The moment of epiphany measure lists the number of subjects who reach epiphany in that particular round as a proportion of all subjects playing in that treatment.

G_{21} -then- G_6 treatment. If the main lesson from G_6 was to mimic the choice of multiples-of-3 rather than transferring the knowledge that G_{21} is of an analytic nature (i.e., achieving mini-epiphany after playing G_6) we might have expected this difference in proportions between treatments to be larger. As it is, it is not statistically significant ($z = 1.43$).

4.2. The impact of playing a simple game first

In light of our support for PR(i) and PR(ii), we now proceed to consider our first main research hypothesis: Subjects playing G_6 figure out that an analytic solution is possible. It dawns on them that there may be an optimal way to play G_{21} too, or that mimicking the strategy of playing multiples-of-3 may work in G_{21} . Even if they do not figure out the optimal strategy right away, on balance they will play G_{21} better in the G_6 -then- G_{21} than in the G_{21} -then- G_6 treatment.

We approach this in a few complementary ways. First we ask: does playing G_6 before G_{21} facilitate complete epiphany? Recall from the introduction our terminology that a subject has reached complete epiphany if he discovers and understands the dominant strategy in G_{21} . This is a cognitive concept which we can only study indirectly. We compare frequencies of perfect play by *Green* players in G_{21} across treatments. The idea is that if a subject plays G_{21} perfectly this probably was no fluke; it is an indicator of complete epiphany. Table 3 records the relevant data. *Green* players play G_{21} perfectly in the G_6 -then- G_{21} treatment 37% of the time, compared to 21 percent in the G_{21} -then- G_6 treatment. This difference is significant at the 5% level (Z statistic = 2.20).

The perfect play test refers to *Green* players only, and neither considers *White* players nor the dynamics as a subject plays five rounds of G_{21} . We next introduce a more dynamic new metric which considers all players: *moment of epiphany*. To help with the definition, we first introduce a notion of *indicated rationality*.

Definition 1. Consider i 's choice $x < 21$ in a given instance of G_{21} . This choice exhibits *indicated rationality* if x is the smallest number such that: it is a multiple-of-three and all i 's subsequent choices in that game are all multiples-of-three.

Although it is technically possible for a winner's indicated rational choice to be 21, she would be a very lucky winner and 'rationality' would play little part. It would require a glaring failure of rational play by at least one party for someone to pass up 18 or fail to win from 18.

Definition 2. Subject i 's *moment of epiphany* occurs in round $R \in \{1, 2, 3, 4, 5\}$ if R is the earliest round such that i has a choice x that exhibits indicated rationality, and in any rounds $R' > R$ all i 's choices exhibit indicated rationality at the earliest available multiple-of-three. We additionally impose that i has a *moment of epiphany* only if she plays at least one round perfectly. (Note that a subject i may have no moment of epiphany.)

Indicated rationality attempts to capture the number of steps-of-reasoning a subject displays in a particular game, while *moment of epiphany* tries to capture the moment when a subject works out the dominant strategy, i.e., when she achieves complete epiphany. Both measures are imperfect. For example, a player may stumble onto the choices 15, 18, and 21 without having achieved epiphany, and yet we would record 15 as the indicated rational choice.¹⁹ Some imprecision seems unavoidable in any measure. We would be wary especially when using these notions to obtain measures of any individual's degree of rationality or insight.

¹⁹ Similarly, the following example shows that there is imprecision associated also with Definition 2. Suppose a player in the *White* position has fully worked out the dominant strategy of playing multiples-of-three, but faces a *Green* player in that round who plays the dominant strategy, picking multiples-of-three at each turn. Then the *White* player is denied an opportunity to pick a multiple-of-three in that round, and will have his moment of epiphany delayed by this measure.




Fig. 1. Histogram of indicated rationality.

In Definition 2, we impose the restriction that a player has a moment of epiphany only if s/he plays at least one round perfectly; the rationale is that this seems more defensible than the assertion that someone who still made mistakes in the final round had nonetheless understood the dominant strategy.

Looking at play in G_{21} across our two treatments, we first note that the median choice with indicated rationality in the G_6 -then- G_{21} treatment is 6, while the median choice with indicated rationality in the G_{21} -then- G_6 treatment is 12 (see Table 3 and Fig. 1). In a sense, the median number of steps-of-reasoning in the G_{21} games in the G_{21} -then- G_6 treatment is three (the steps involving 12, 15, 18, and 21; cf. footnote 4), while the median number of steps-of-reasoning in the G_{21} games in the G_6 -then- G_{21} treatment is five (the steps being 6, 9, 12, 15, 18, and 21). Thus, playing G_6 prior to playing G_{21} seems to increase the median steps-of-reasoning achieved by subjects from three to five in G_{21} . A two-sample Kolmogorov–Smirnov test for equality of distributions shows we can reject the null hypothesis of both treatments being drawn from the same distribution at the 5% level ($p = 0.025$). A k -sample test on the equality of medians yields the same conclusion.

We next look at the distribution of the subjects' moment of epiphany across the treatments. An important statistic here is the number of players who never achieve epiphany: 53 percent of players in the G_6 -then- G_{21} treatment and 64 percent in the G_{21} -then- G_6 treatment (see Table 3). These proportions are not significantly different ($Z = 0.93$) from each other, and seem to suggest that some players may never work out the dominant strategy in G_{21} ; there just may be too many steps-of-reasoning involved. This is despite many of these subjects successfully working out the dominant strategy in G_6 (either before or after playing G_{21}).²⁰

In addition, the data suggests that playing G_6 before playing G_{21} may not help this group of subjects achieve epiphany. However, among the group of players who do achieve epiphany (as indicated by our measure based on Definition 2), it appears that playing G_6 first does help some subjects achieve complete epiphany sooner. In the G_6 -then- G_{21} treatment, up to 37 percent of players (11/30) achieve their moment of epiphany by round 3, as opposed to 19 percent of players (8/42) who achieve their moment of epiphany by round 3 in the G_{21} -then- G_6 treatment, a statistically significant difference ($Z = 1.67$). More generally, the two-sample Kolmogorov–Smirnov test for equality of distributions also shows we can reject the null hypothesis of both treatments being drawn from the same distribution at the 5% level ($p = 0.017$) and the k -sample test on the equality of medians confirms this conclusion. A kernel density plot across 5 rounds of play (see below for explanation) illustrates this difference in Fig. 2.

We noted earlier that the moment of epiphany measure has some short-comings (e.g. footnote 13). While it shows that some particular player understood the dominant strategy by some moment, it does not exclude the possibility that it was understood *before* that moment. To view the data from yet another angle, we define *epiphany delay*, which identifies the *latest* moment when complete epiphany was demonstrably *not* achieved. To this end, we generate a scale $0 \rightarrow 105$ as follows: Each position in the first round of G_{21} is assigned $1 \rightarrow 21$; each position in round 2 is assigned $22 \rightarrow 42$ and so on up to $85 \rightarrow 105$ for the final round. We use this scale to record the *last* occasion when a player *failed* to choose an available multiple-of-three.

Definition 3. Subject i 's *epiphany delay* is an element of $\{0, 1, \dots, 105\}$ identifying the *last* occasion when i fails to select a multiple-of-three when able to do so.²¹ If i never misses such an opportunity, then we assign 0 as i 's epiphany delay measure.

²⁰ Once again this seems to point to the second learning mechanism – involving mini and complete epiphanies, rather than to the 'mimicking' mechanism. If players were simply mimicking the strategy they used in G_6 then we would expect much higher numbers of players in the G_6 -then- G_{21} treatment to achieve complete epiphany in the G_{21} game.

²¹ Subject i is able to choose a multiple-of-three every time his/her opponent has not played a multiple-of-three in his/her turn.




Fig. 2. Kernel density of epiphany in two treatments.

Comparing epiphany delay in G_{21} across the two treatments, we find that in the G_6 -then- G_{21} treatment the mean epiphany delay is 53.66 with a standard deviation of 36.53. In G_{21} -then- G_6 the mean epiphany delay is 68.93 with a standard deviation of 29.2. That is, on average the location of the last error subjects make in the G_6 -then- G_{21} treatment is 10 or 11 in the third round, while in the G_{21} -then- G_6 treatment it is 6 in the fourth round. The difference between the means across these two treatments is significant ($Z = 1.90$) at the 5% level.

A two-sample Kolmogorov–Smirnov test for equality of distributions narrowly fails to reject the null hypothesis of both treatments being drawn from the same distribution at the 10% level ($p = 0.102$) although a k -sample equality of medians test rejects the null at the 10% level ($p = 0.056$). The above kernel density plot in Fig. 3 is also suggestive of a difference between treatments.

Taking these various measures together, we conclude that although one or two achieve only marginal statistical significance, *all* the differences are in the predicted direction and *most* are strongly so. This suggests that prior experience with a simple game of suitable structure does indeed induce learning, which then raises the likelihood of, or advances the moment when, epiphany will be achieved in a similar game of greater depth.




Fig. 3. Kernel density of epiphany delay.

Table 4
Median indicated rationality.

	G_6 -then- G_{21}
Round 1	15
Round 2	12
Round 3	6
Round 4	4.5
Round 5	3

4.3. Learning within G_{21}

We have seen that it takes time for subjects to learn to play the dominant strategy in G_{21} . In this section we focus on *how* learning happens in G_{21} given that subjects have played the simpler G_6 game before. Thus, we focus on learning patterns from the G_{21} games in the G_6 -then- G_{21} treatment.²²

To consider learning, let us first look at how the indicated rationality measure evolves across rounds (see Table 4).

After playing 5 rounds of G_6 , the median indicated rationality of subjects in round one is 15. Interpretation: after realizing that an analytic solution is possible, at least 50 percent of subjects appear to work out two steps-of-reasoning (cf. footnote 3) in the first round of G_{21} .²³ In round two, they seem to work out one more step, with the median indicated rationality falling to 12. Then there is a jump, with median indicated rationality falling to 6 by round three, indicating five steps-of-reasoning by at least half of the subjects playing in this round. The median indicated rationality of 3 by round five reflects the fact that, by the beginning of that round, almost half of the subjects have reached complete epiphany, and have worked out the dominant strategy. It is interesting to note that two levels of reasoning would seem to lead to insight on how to play G_6 perfectly, and that the median subject playing G_{21} for the first time also seems to be able to reason out two steps. This may explain why although so many subjects (93 percent across the two treatments) play perfectly in G_6 , many of the same subjects fail to achieve complete epiphany in G_{21} even after five rounds.

While there appears to be a steady learning process as the rounds progress on the part of the median subject, there is a substantial degree of heterogeneity in how quickly subjects learn to play the dominant strategy.²⁴ Many seem never to learn how to play (as indicated by their moment of epiphany measure), while a number of them appear to reach complete epiphany at or near the beginning.²⁵ So at any given point during play of a series of G_{21} there is likely to be (i) players who have reached complete epiphany; (ii) other players who have understood that an analytical solution is possible (or have otherwise gained some insight) through playing G_6 first and are lowering their onset of indicated rationality in subsequent rounds; and (iii) yet other players who have failed to achieve any insight and who still ascribe victory to chance.

We note an interesting difference between the two treatments in terms of learning when we confine ourselves just to those subjects who *fail* to reach complete epiphany. It appears that in the G_{21} -then- G_6 treatment, these subjects nevertheless are making progress towards learning the dominant strategy; in round one their median indicated rationality is 18 and by round five it falls to 10.5. On the other hand, in the G_6 -then- G_{21} treatment there is no evidence of learning by this (smaller) group of players; in round one their median indicated rationality is 15 and by round five it remains at 15.

5. Conclusions

How do you defeat a Gordian knot? How do you make an egg stand on end? Wise men failed to come up with answers until Alexander the Great and Christopher Columbus came along. Their legends teach us about how fame and fortune may be the reward for clever insights.

How do humans play games? We suggest that an adequate answer requires understanding *how* humans reach clever insights. Most of economic theory assumes that decision-makers best respond to their beliefs. Yet optimizing is often complicated and there is an abundance of related issues to explore: Do decision-makers understand when problems admit analytical solutions? Does the answer to the previous question depend on their life experiences? How efficient are humans in calculating solutions? What are the processes by which they learn to optimize?

²² Note that using data from the G_{21} -then- G_6 would be inappropriate for answering the questions in this section since whatever delay in correct choices occurs may depend on the absence of understanding that an analytic solution is possible in that treatment. However, we sometimes report the results from the G_{21} -then- G_6 treatment, mainly to contrast it with the G_6 -then- G_{21} treatment.

²³ Once again it is important to note here that, had the learning mechanism from the simpler G_6 game to the more complicated G_{21} game been 'mimicking,' one would likely observe much lower median indicated rationality in the earlier rounds of the G_{21} game in the G_6 -then- G_{21} treatment. The fact that is not the case, and that median indicated rationality falls across rounds (as seen in Table 4) is another piece of evidence that the learning mechanism is likely through the mini-epiphany and then complete epiphany mechanism.

²⁴ This result joins a wealth of research suggesting ways individuals differ, from the 'Big five' personality dimensions of extraversion, agreeableness, conscientiousness, neuroticism and openness to experience (John and Srivastava, 1999) to psychological inclinations like sensitivity to emotional concerns (e.g. Krone, 2003) or (of more relevance to us) level of thinking regarding the rationality and beliefs of others (e.g. many of the references cited in Section 2).

²⁵ Thirty percent of subjects have their moment of epiphany by round 2 (see Table 3).

We explore related issues in connection with a two-player zero-sum game of perfect information: G_{21} is much simpler than chess, possible to figure out optimal play for, and yet sufficiently complicated that most humans do not figure it out at least at first. The borderline solvability-by-humans makes it suitable as a research tool for shedding light on questions like those in the previous paragraph.

To structure our examination of human insights in games, we introduce a second, but shorter game, G_6 , and study how playing G_6 first facilitates learning to play optimally in G_{21} . There are two plausible learning mechanisms—(a) mimicking or (b) learning that an analytic solution is possible in the shorter game and transferring this knowledge to the more complicated game.

We propose an experimental design to study how learning transfer occurs across the games of G_6 and G_{21} . Our results confirm that experience with G_6 improves performance in G_{21} (according to a variety of measures). Our evidence also seems to point to the fact that learning across these two games occurs through the second mechanism (achieving a mini-epiphany that an analytical solution is possible and then attaining nirvana (complete epiphany) when the dominant strategy of the more complicated game is learned). If learning were mainly through mimicking we would have been likely to observe a lot more perfect play as well as much earlier moments of epiphany in the G_{21} game in the G_6 -then- G_{21} treatment. Instead, what we observe is a more gradual unwrapping of the optimal strategy: consistent with the player having realized an analytic solution is possible, but working this analytic solution out in steps.

We also examine the nature of learning once the subject realizes that an analytic solution is possible. Here we do not have a preconceived hypothesis. It turns out that learning towards complete epiphany is gradual to some degree in most subjects. However, subjects exhibit a lot of associated heterogeneity. Experience matters in possibly predictable ways, but there is a lot of individual variation.

Very little discourse in economics seems to be concerned with how human minds get primed to engage in rational thinking, and how insights are reached.²⁶ More research is concerned with how players reason about others (see our Section 2 and the references to the literature on level- k and cognitive hierarchy models). We suggest that these research goals are complementary, and that future work should keep both goals in mind.

While humans may have a language instinct with which to acquire proficiency in spoken language, strategic thinking, like written language, has to be learned the hard way. The connections between our findings and broader questions, such as why societies value schools, or how we may best structure teaching to foster insight and improve learning (e.g., begin with the simplest example of a concept), should be kept in mind although at the moment tackling such questions remains beyond our scope. Applications however could include the facilitation of learning how to plan ahead when actions are needed today but the consequences are temporally distant. We have not attempted here to examine particular economic settings but one may suspect that the array of relevant ones is vast and could include, for example, choices regarding education, nuclear waste storage, marketing, and warfare.

Acknowledgments

We would like to thank Pierpaolo Battigalli, Harry Bloch, Alex Coram, Ken Clements, Vince Crawford, Brit Grosskopf, Willemien Kets, Paul Miller, two referees, a Co-Editor, and several seminar attendees for their comments. We also appreciate the excellent research assistance that Dahai Fu provided.

Appendix A.

Instructions, game sheets, guide cards for game matchings:

{Subjects' instructions were written on the same page as the game sheets. We explained verbally that movements for game matching should proceed according to schedule cards that we distributed. These assigned students to "tables" and explained who would act as the White/Green player. We indicate here the look of the instruction/game sheet for the game of twenty-one (the game of six was handled analogously) and the schedule card for one of a group's subjects.}

A.1. The game of 21

Welcome! The rules of the game are:

Each player takes turns playing the game, with the white player beginning. To begin, white can choose either the number 1 or the number 2, by circling one of them. The green player then plays by incrementing white's choice by 1 or by 2. That is, if white had circled the number 1, then green can choose either the number 2 or the number 3. If, instead, white had chosen

²⁶ Although not game-theoretically anchored, discussions of entrepreneurial discovery and creativity have some of this flavour. See e.g. Hayek (1978/1984) and Kirzner (1985) for arm-chair reasoning, and Demmert and Klein (2003) for a related out-door experiment. They test whether the strength of financial incentives matters to whether subjects figure out the most efficient method (inverting of a plastic stool) for transferring water from spot A to spot B, thereby getting insights-by-analogy on a conjecture about entrepreneurial discovery by Hayek and Kirzner. Demmert and Klein use epiphany to refer to entrepreneurial discovery but their usage differs from ours as they exclude understanding that results from deliberate effort.

the number 2, then green can choose either the number 3 or the number 4. Green uses a cross to mark his/her choice. The game continues with each player incrementing the other's choice by 1 or by 2, until one player reaches 21. The player who reaches 21 first wins.

WHITE PLAYER: Circle the number you choose in each round.

GREEN PLAYER: Use a cross to mark the number you choose in every round.

WHITE ALWAYS BEGINS. PLAYER WHO REACHES 21 FIRST WINS

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
---	---	---	---	---	---	---	---	---	----	----	----	----	----	----	----	----	----	----	----	----

WHITE PLAYER ID: _____

GREEN PLAYER ID: _____

WINNER ID: _____

PLAYER B's SCHEDULE

GAME 1:	Plays against A at Table 1.	Position: Green
GAME 2:	Plays against C at Table 3.	Position: White
GAME 3:	Plays against E at Table 2.	Position: Green
GAME 4:	Plays against F at Table 2.	Position: White
GAME 5:	Plays against D at Table 3.	Position: Green

References

Bolle, F., 1990. High reward experiments without high expenditure for the experimenter? *Journal of Economic Psychology* 11, 157–167.

Bosch-Domenech, A., Montalvo, J., Nagel, R., Satorra, A., 2002. One, two (three), infinity...: newspaper and lab beauty-contest experiments. *American Economic Review* 92, 1687–1701.

Bouton, C.L., 1901–02. Nim, a game with a complete mathematical theory. *Annals of Mathematics* 3, 35–39.

Camerer, C., 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, New Jersey.

Camerer, C., Ho, T.-H., Chong, J.-K., 2004. A cognitive hierarchy model of games. *Quarterly Journal of Economics* 119, 861–898.

Chou, E., McConnell, M., Nagel, R., Plott, C., 2009. The control of game form recognition in experiments: understanding dominant strategy failures in a simple two-person guessing game. *Experimental Economics* 12, 159–179.

Cooper, D., Kagel, J., 2008. Learning and transfer in signaling games. *Economic Theory* 34 (3), 415–439.

Cooper, D., Kagel, J., 2009. The role of context and team play in cross-game learning. *Journal of the European Economic Association* 7, 1101–1139.

Costa-Gomes, M., Crawford, V., 2006. Cognition and behavior in two-person guessing games: an experimental study. *American Economic Review* 96, 1737–1768.

Costa-Gomes, M., Crawford, V., Broseta, B., 2001. Cognition and behavior in normal-form games: an experimental study. *Econometrica* 69, 1193–1235.

Crawford, V., 2003. Lying for strategic advantage: rational and boundedly rational misrepresentation of intentions. *American Economic Review* 93, 133–149.

Crawford, V., Gneezy, U., Rottenstreich, Y., 2008. The power of focal points is limited: even minute payoff asymmetry may yield large coordination failures. *American Economic Review* 98, 1443–1458.

Crawford, V., Iriberry, N., 2007a. Fatal attraction: salience, naivete, and sophistication in experimental 'hide-and-seek' games. *American Economic Review* 97, 1731–1750.

Crawford, V., Iriberry, N., 2007b. Level-k auctions: can a non-equilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica* 75, 1721–1770.

Demmert, H., Klein, D., 2003. Experiment on entrepreneurial discovery: an attempt to demonstrate the conjecture of Hayek and Kirzner. *Journal of Economic Behavior and Organization* 50, 295–310.

Dixit, A., 2005. Restoring fun to game theory. *Journal of Economic Education* 36, 205–219.

Ewerhart, C., 2000. Chess-like games are dominance solvable in at most two steps. *Games and Economic Behavior* 33, 41–47.

Fudenberg, D., Kreps, D., 1988. *A Theory of Learning, Experimentation and Equilibrium in Games*. MIT, Mimeo.

Gabaix, X., Laibson, D., 2000. A boundedly rational decision algorithm. *American Economic Review: Papers and Proceedings* 90, 433–438.

Gneezy, U., 2005. Step-level reasoning and bidding in auctions. *Management Science* 51, 1633–1642.

Gneezy, U., Rustichini, A., Vostroknutov, A., this issue. Experience and insight in the race game. *Journal of Economic Behavior and Organization*.

Grosskopf, B., Nagel, R., 2008. The two-person beauty contest. *Games and Economic Behavior* 62, 93–99.

Hayek, 1978/1984. Competition as a discovery procedure. Reprinted in: *The Essence of Hayek*. Hoover Institution Press, Stanford.

Ho, T.-H., Camerer, C., Weigelt, K., 1998. Iterated dominance and iterated best response in experimental 'p-beauty contests'. *American Economic Review* 88, 947–969.

John, O., Srivastava, S., 1999. The big five trait taxonomy: history, measurement, and theoretical perspectives. In: Pervin, L., John, O. (Eds.), *Handbook of Personality: Theory and Research*, 2nd ed. Guilford, New York.

Johnson, E., Camerer, C., Sen, S., Rymon, T., 2002. Detecting failures of backward induction: monitoring information search in sequential bargaining. *Journal of Economic Theory* 104, 16–47.

Kirzner, I., 1985. *Discovery and the Capitalist Process*. University of Chicago Press.

Krone, H.W., 2003. Individual differences in emotional reactions and coping. In: Davidson, R., Scherer, K., Goldsmith, H. (Eds.), *Handbook of Affective Sciences*. Oxford University Press, New York.

McKinney, N., Van Huyck, J., 2006. Does seeing deeper into a game increase ones' chances of winning? *Experimental Economics* 9, 297–303.

- McKinney, N., Van Huyck, J., 2007. Estimating bounded rationality and pricing performance uncertainty. *Journal of Economic Behavior and Organization* 62, 625–639.
- Nagel, R., 1995. Unravelling in guessing games: an experimental study. *American Economic Review* 85, 1313–1326.
- Östling, R., Wang, J., Chou, E., Camerer, C., 2009. Testing game theory in the field: Swedish LUPI lottery games. SSE/EFI Working Paper 671, Stockholm School of Economics.
- Selten, R., Abbink, K., Buchta, J., Sadrieh, A., 2003. How to play (3×3)-games: a strategy method experiment. *Games and Economic Behavior* 45, 19–37.
- Stahl, D., Wilson, P., 1994. Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization* 25, 309–327.
- Stahl, D., Wilson, P., 1995. On players' models of other players: theory and experimental evidence. *Games and Economic Behavior* 10, 218–254.