Skill in Games

Patrick Larkey • Joseph B. Kadane • Robert Austin • Shmuel Zamir

H. J. Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213
Department of Statistics, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213
Harvard Business School, Cambridge, Massachusetts 02139
Department of Statistics, Hebrew University, Jerusalem, Israel

Differences in players' skill are important determinants of relative player success in most *real games* such as poker, chess, basketball, business, and politics. Yet conventional game theory has concentrated primarily on games with no skill differences among players. This paper uses a simplified version of stud poker to better understand the concept of differential player skill in games. Players with very different strategies for playing this game are modeled algorithmically and pitted against one another in simulation tournaments.
(*Bayesian Game Theory; Game Theory; Poker; Simulation; Strategy*)

1. Introduction

Skill is the extent to which a player, properly motivated, can perform the mandated cognitive and/or physical behaviors for success in a specific game. More skillful players tend to score better than less skillful players. The skills essential to success tend to be game specific; success in basketball requires a somewhat different set of skills than success in chess, business, or politics.

Skill differences among players are important to most real games' outcomes. Clearly, theories that aspire to predict outcomes or to advise players on how to improve their play—two of the more obviously useful potential "applications" of a theory of games—require some explicit representation of differential player skills. The mathematical theory of games as it has evolved over nearly a half century since the seminal work of John von Neumann and Oskar Morgenstern (1944) has not, however, with a few exceptions, explicitly incorporated skill differences among players as an important feature of the games studied. Conventional theories of games focus on equilibria and other "solution concepts" resulting from unboundedly rational players utilizing optimal or equilibrium strategies. Differences in game structures and in player endowments in information and resources have been featured in theories of games.

The theory of games has been primarily concerned with limiting cases in which skill differences among players are not important. The theoretical assumption about skill, usually implicit, is that all players have incentives to play optimally and that they will figure out how to do so, at least in the long run. The cognitive or physical difficulties for players in devising and executing strategies for playing particular games are essentially assumed away.

While there is venerable work on the topic of boundedly rational game players (Simon 1957), the issue has only recently garnered much interest in modern decision theory, game theory, and economics; Kalai (1990) reviews a growing body of work on bounded rationality, memory, and complexity. Camerer (1990) explores the basis for a "behavioral game theory" in more descriptively adequate assumptions about player behavior. Binmore (1990) provides a sweeping, critical review of game theory and concludes that "an attempt must be made to model players' thinking processes explicitly." Binmore echoes Herbert Simon who has been providing procedural representations of human reasoning and arguing cogently on their behalf for forty years (Simon 1955, 1959, 1976, 1983, and 1991). Axelrod and Hamilton (1981) and Axelrod (1984) found important skill differences among players in a computer-based Prisoners' Dilemma tournament. Leifer (1991) analyzes chess and the relationship between skill and the

social relationships among players. Beasley (1990) explores the mathematics of games of varying skill.

This paper seeks to understand the concept of skill in games as an initial step toward building theories of real games capable of both predicting outcomes and advising play. The analytic core of the paper is a detailed analysis of a game of skill, Sum Poker. We posit players in the form of alternative strategies for play that differ in the information they use and in how they use it. The general behaviors mandated for player success in Sum Poker—observation, memory, computation, knowledge of the random device, misleading opponents about the actual strength of your position, and correct interpretation and forecasts of opponents' behaviors—are common to many "real games." The levels of skill and relative success of the different strategies are explored in computerized experiments.

2. A Skill Typology of Games

Game theorists categorize games on several different dimensions. Games are cooperative or noncooperative and one-stage or multi-stage; games have payoffs that are zero-sum or nonzero-sum and involve 2 to n players; the information available to players in a game is complete or incomplete and then, symmetric or asymmetric. When experts on game playing and gambling (Scarne 1980, Jacoby 1963, Thackrey 1971, Morehead 1967, and Livingston 1971) categorize games, they completely ignore the usual game-theoretic dimensions and focus on another dimension, **skill**. The critical aspect of games for these experts is the extent to which outcomes depend on player skill rather than luck. Their books are primarily about acquiring skill. Another significant difference between game theorists and game-playing experts is in the scope of what constitutes a game. Game theorists' restrictively define a **game** as the full description of its rules. The game playing experts include in their definition of game, always implicit and informal in their discussions, many factors beyond the formal rules including the decisions on which games to play and information on the past performance of opponents. The experts' notion is roughly equivalent to the game theorists' notion of a **repeated game** or **supergame** but with an ill-defined *beginning* and *horizon*.

There are three types of games in terms of skill: (1) In **Pure Chance** games such as lotteries, Keno, Matching

Pennies, War, and Show Down Poker players compete against a random device that cannot be influenced; the probability of a particular player winning is simply a function of the game's random device. Other games such as Roulette, Craps, and Chuk-a-Luck are essentially Pure Chance games but often have wagering rules appended that introduce elements of skill; more skillful players can lose their money at slower rates; (2) **Skill-Chance** games such as Poker, Backgammon, Rubber Bridge, and Gin Rummy have both a random device—cards or dice—and significant elements of skill (see Kadane (1986) on electronic draw poker); (3) **Pure Skill** games such as Tic-Tac-Toe, Go Moku, Checkers, and Chess have no external chance elements. The probability of a particular player winning is essentially a function of that player's skill relative to the other players.

2.1. The Skill Concept

Players face three critical actions with respect to games of all kinds. First, they can usually choose the games in which they will participate; there are usually many more games available than a player has the resources, time, and money, to play. Second, given a decision to play, players must plan a strategy for play. Third, given the decision to play and a strategy, players must execute their planned strategy.

Each of these three actions entails a different type of skill: (a) **Strategic Skill** is the ability to decide what games to play. The boundedly rational player selects games in which his skills (i.e., those associated with choosing and executing strategies during a game) relative to other players' skills yield positive expected utility, including utility from both the process of playing games and expected payoffs. The more accurate the judgments about relative levels of game skill, the greater the strategic skill. Strategic skill is roughly analogous to skill in the play of a *supergame*; (b) **Planning Skill** is the ability to formulate strategies relative to specific opposition in a specific game. Formulating strategies is largely the cognitive activity of creating a planned course of action for conduct in the game. Two key components of planning skill are self and opposition assessment; (c) **Execution Skill** is the ability to execute a planned strategy. It may be much easier to conceive a strategy that entails remembering all of the cards played and a detailed record of opponents' past behaviors to

be used in predicting their future behaviors than it is to accomplish these feats of memory (and subsequent analysis) in "real time" with no external memory aids.

In games with a random device (e.g., dice, cards, economic climate, etc.) these three skills only partially determine levels of success in a particular game. The relative importance of each action and of the random device in determining outcomes depends on the details of the game. These skills interact in their effects. Choosing a game in which you are overmatched and, regardless of how well you strategize and play, you will probably lose. Choose a game where you are undermatched, can make victory likely, if not certain. Assessing your skills relative to prospective opponents in choosing games to play usually involves some planning skills; you must imagine your strategy and predict both the opponent's strategy and the likely outcomes to do the assessment. Many games are played with little or no conscious planning; the strategy is implicit and adaptive in the play.

2.2. Representing Skill

There are two fundamentally different ways of representing the skill of a particular individual with respect to a particular game. **Process representations** describe players by their method of play. For poker, the method of play includes the informational and decisional procedures—what is noticed and remembered and the set of rules used to choose an action at each juncture in the game. For example, two players might be identical in all process respects except one remembers all of the cards that have been played while the other forgets about half of them. **Performance representations** characterize players by outcomes. For example, in a regular weekly six-hand poker game one player is $10,000 behind and another player is $10,000 ahead over the last year. In real situations there is usually some information on both types of representation. Process representations tend to be more useful in advising play, while performance representations are more useful in predicting game outcomes.

Both types of skill representation have problems as theoretical tools for the study and play of games. **Process representations** are positive theories of player behavior that can, in principle, provide normative information about playing by comparing more- and less-successful players utilizing different methods. Coherent,

valid, normative theories of playing methods that lead to success, much less optima, in most real games rarely exist because: (1) the methods of great players are hard to describe and emulate; (2) the strategy space for real games tends to be too large and complex to describe; and (3) optimal play can only be defined conditional on the characteristics of opponents. Extant theories of play tend to be neither necessary nor sufficient to success. A poker player who forgets about half of the cards that have been played should be less successful over time than an otherwise comparable player who remembers them all. But as the analysis of simulated games below shows, faulty memory can be indistinguishable from bluffing, and players with this flaw may defeat a nonbluffing player who remembers everything.

Players with the higher frequency of successful outcomes in an adequate sample of play are, ceteris paribus, more skillful. Unfortunately, ceteris non paribus in virtually all real games and performance representations are rarely straightforward. One pandemic problem is that play outcomes only have meaning relative to the competition. For example, losing a game of chess in 20 moves to Gary Kasparov is surely a much more impressive performance accomplishment than winning a game in 20 moves over an eight-year-old who has just learned what constitutes a legal move for each piece. But how much more? To construct a proper performance record for a player requires adjustment for the strength of the competition, but there is no demonstrably correct method of adjustment because there is no ultimate source of correct skill rankings that can be used to assess the adjustment method.[1] Sampling also becomes an issue because most real games are skill-chance games and observed outcomes confound skill and chance. In principle, skill and chance can be separated statistically with an appropriate set of experiments, viz., identical, independent repetitions. In practice, the experiments arising in the course of play of real games are far from ideal and pose significant challenges to inferential learning.

---

[1] See Larkey (1991) for adjustments to compare professional golfers; Caulkins et al. (1993) for adjustments for comparing airlines' on-time performance; and Larkey and Caulkins (1992) for adjustments to grade point averages.

3. Sum Poker

Sum poker, a simplified form of stud poker, preserves the essential elements of poker—memory, bluffing, and observing the play of others. The simplifications, fewer cards and a less complicated ranking of hands, are insignificant departures from real poker. Sum poker has the following sequence of play: (1) everyone antes the amount of the minimum bet; (2) two cards are dealt to each player, one down and one up; (3) a round of betting ensues; (4) another card is dealt up; (5) another round of betting follows; and (6) hands are scored by the sum of the three cards; face cards count ten, aces one, and all other cards receive face value. The highest sum wins.

During a betting round, a player may, depending on position, open, raise, call, fold, or check. To open means to initiate the betting round with a positive bet. To raise means to bet an amount above a pending bet from another player; other players are obliged to either match the bet or concede the hand. To call means to bet exactly the amount of a pending bet from another player (or players in the case of a raise), thereby staying in the hand. To fold is to concede the hand to the opponent. To check—remaining in the hand without betting or folding—is an option only when there is no prior bet in the round.

The opening bettor in each round is chosen at random. No more than three raises are allowed in a betting round. There is a maximum and minimum betting amount, $5 and $10 respectively, unless otherwise indicated. The game is played repeatedly, in some cases tens of thousands of hands, until only one player remains who is not ruined, viz., out of money.

3.1. Players

In our game simulations, we posit and test 12 different strategies for playing Sum Poker. The 12 players briefly described[2] are:

1. **Simple [S]** is a baseline strategy that plays randomly without regard to any of the possibly relevant information about the deck, own hand, or opponents'

cards and behavior. Beyond knowing the legal moves in the game, this player arguably has no skill.

2. **AvgHand [AH]** compares cards with an average hand [13.08 in the first round (two cards) and 19.62 in the second round (three cards)] and bets accordingly. Where **AvgHand's** cards are much better than average, she bets very aggressively. **AvgHand** completely ignores opponents' cards.

3, 4, and 5. **Loose, Middle, and Tight [L, M, T]** are three closely related strategies. They make different assumptions about the opponent's down card. **Loose** is very aggressive and always assumes that the opponent's down card is a 1. **Middle** assumes a 5. **Tight** assumes a 10. The strategies then consist of rules conditioned on a comparison of own and opponent's hands given the assumption.

6. **BluffsLots [BL]** bluffs a lot. Any time that **BluffsLots'** up cards are greater than the opponents' up cards or a simple calculation of probable advantage is in his favor, he raises the maximum amount. **BluffsLots** uses information about the deck, about cards played, about his own hand, and about his opponents' hands.

7. **CalcMuch [CM]** is a nonbluffing strategy that does not learn about the opponent and uses a relatively simple set of calculations. The strategy uses knowledge of what cards have been played to compute the probability of winning in the cards against a specific opponent. There is no consideration of money.

8. **PlayerCalc [PC]** is a learning version of **CalcMuch**. The strategy observes and remembers for each opponent for each possible difference in up cards what the lowest value of a down card has been that the opponent has ever been willing to bet on. There is no bluffing and no consideration of money. **PlayerCalc** reverts to **CalcMuch** when an opponent's bluffing successfully counters its learning mechanism.

9. **MixedCalc [MC]** is a *complex strategy* that plays as **CalcMuch** with probability 0.94 and as **Simple** with probability 0.06.

10. **PlayerCalcB [PCB]** is a complex strategy which plays as **PlayerCalc** with probability 0.94 and as **Simple** with probability 0.06.

11. **ExpVal [EV]** computes expected values and plays accordingly. The probability portion of the expectation is the same as **PlayerCalc**. The value portion is based upon the difference between an estimate of what

---

[2] The appendix provides a fuller technical description of the conditions and actions constituting each strategy. The PASCAL code is available from the authors.

the final pot will be if neither player folds and an estimated cost of staying in the hand. There is no bluffing.

12. **ExpValB [EVB]** is a bluffing version of **ExpVal**. It plays as **ExpVal** with probability 0.94 and as **Simple** with probability 0.06.

3.2. Process Representation of Strategies

Table 1 shows process characteristics of the strategies on several dimensions. The **knowledge** dimension is the amount of information in the game environment that the strategy can use. **Computation** summarizes the computational ability required to execute the most complex rules in the rule set. **Algorithm length** is measured lines of code as a rough proxy for the size of rule sets. **Conservatism** refers to how certain of victory a strategy must be to bet. **Learning** and **Bluffing** simply indicate whether those abilities are built into the strategy intentionally.

The *strategy space* for playing Sum Poker is huge. We cannot know analytically that our twelve strategies include the strategy that will perform best on average against all other possible strategies or best against any particular strategy. The strategies used in the simulations were not, however, posited randomly, but systematically cover: (1) the potential information available to play (from no information to full—including all past—information about the deck, current cards, and opponents); (2) the potential modes of computation (from none to computations well beyond human capacities); (3) the risk propensity (from very aggressive to very cautious); and (4) the extent of attempts to conceal patterns of play from opponents. The purpose of these strategies is to explore the concept of skill in games, not to simulate the behavior of actual subjects playing sum poker.

3.3. Performance Representation of Strategies

Each strategy was matched against each of the others in a simulated tournament. A game in a tournament consisted of the hands sufficient for one player to ruin the opponent. Each pair of players was simulated in 100 games.

Table 2 summarizes the results of the tournament. The number in each cell tells the number of times in 100 games that the player at the left beat the player at the top. The Total column contains a simple sum of the

Table 1   Characteristics of Strategies

	Knowledge	Computation	Algorithm Length	Conservatism	Learning	Bluffing
Simple [S]	0	0	0.43	2	No	No
AvgHand [AH]	1	1	0.61	2	No	No
Loose [L]	2	2	0.79	1	No	No
Middle [M]	2	2	0.96	2	No	No
Tight [T]	2	2	0.96	3	No	No
BluffsLots [BL]	3	3	0.87	1	No	Yes
CalcMuch [CM]	3	3	1.54	2	No	No
MixedCalc [MC]	3	3	1.96	2	No	Yes
PlayerCalc [PC]	4	3	2.46	2	Yes	No
PlayerCalcB [PCB]	4	3	2.89	2	Yes	Yes
ExpVal [EV]	4	4	3.57	2	Yes	No
ExpValB [EVB]	4	4	4	2	Yes	Yes

Knowledge	Computation	Conservatism
0 = knows nothing	0 = none	
1 = knows own cards	1 = counting and comparing	1 = aggressive
2 = knows others' cards	2 = adding and subtracting	2 = moderate
3 = knows distribution of cards in deck	3 = multiplying and dividing	3 = conservative
4 = knows history of play	4 = beyond human abilities	

Table 2   Tournament Results

	PCB	PC	MC	CM	EV	M	EVB	AH	T	BL	L	S	Total	$'s
PCB		81	48	79	92	91	96	99	99	98	99	100	982	432K
PC	19		61	93	90	100	87	100	100	100	100	100	950	400K
MC	52	39		38	96	92	93	97	100	96	98	100	901	351K
CM	21	7	62		82	100	86	100	100	99	100	100	857	307K
EV	8	10	4	18		34	60	38	100	93	97	100	562	12K
M	9	0	8	0	66		73	58	9	90	98	100	511	−39K
EVB	4	13	7	14	40	27		36	98	83	88	100	510	−40K
AH	1	0	3	0	62	42	64		20	88	98	100	478	−72K
T	1	0	0	0	0	91	2	80		5	100	100	379	−171K
BL	2	0	4	1	7	10	17	12	95		74	99	321	−229K
L	1	0	2	0	3	2	12	2	0	26		97	145	−405K
S	0	0	0	0	0	0	0	0	0	1	3		4	−546K

number of games won by each player during the tournament and the net winnings is the amount won or lost by the strategy during the tournament.

Relating the process and performance representations would be simpler if performance relationships between strategies were transitive. That is, if player *A* is better than *B*, and *B* is better than *C*, then *A* should be better than *C*. Only where performance relationships are transitive can there be a single measure of skill that is predictive of outcomes when any two strategies are matched.

A strict, predictive rank ordering of players by performance is not possible. **Tight** beats **Middle** 91 of 100 times. **Middle** is 90 percent victorious against **BluffsLots**. Yet **BluffsLots** wins against **Tight** 95 percent of the time. **Tight, AvgHand,** and **BluffsLots** form another intransitive trio. The intransitivities are not surprising. There is no theoretical reason to expect the rule sets qua players to interact in ways that yield overall summary performance measures. It is a commonly observed phenomenon in real games that players and teams "match up differently."[3]

Intransitivities are neither numerous nor do they span huge differences in total victories or net winnings. While strategies using similar processes (i.e., number and type of rules, etc.) may perform intransitively, adding a significant degree of capability, assuming it is used

---

[3] See Gardner (1974) for a discussion of the implications of intransitivities in a variety of real games.

appropriately, can make a strategy strictly better than another. For example, giving **CalcMuch** the ability to calculate a probability of winning makes it much better than any of the strategies that do not have this ability. Only learning versions of **CalcMuch—ProbCalc** and **ProbCalcB—beat CalcMuch**.

3.4. Execution Skill—Processes, Capabilities, and Performance

There is an imperfect relationship between performance measured as total or net winnings and the internal characteristics—knowledge, computation, algorithm length, learning ability, bluffing ability—our proxy for capabilities in playing Sum Poker. **Simple,** the least capable strategy is the worst performer and **PlayerCalcB,** one of the most capable, is the strongest performer. A strategy's ability to use information, compute, and adapt to a greater variety of specific situations seem to place a ceiling on how good a player can be. Increasing a strategy's capabilities does not, however, guarantee improved performance. For example, the most "capable" strategies in terms of process, **ExpVal** and **ExpValB,** are not particularly strong performers. Conversely, **CalcMuch,** a nonlearning, nonbluffing strategy with a relatively short algorithm, plays very well.

3.5. Planning Skill—Conservative vs. Aggressive Choices

One important dimension of a strategy is *conservatism*—how aggressively will a strategy bet in particular game situations? An important determinant of *conservatism* is

how optimistic or pessimistic a player is about the value of an opponent's hole card. The most conservative strategy, **Tight,** assumes the worst: the opponent has a ten down. The least conservative (most aggressive) strategy, **Loose,** assumes the best: the opponent has a one down. The other strategies, including **Middle,** which assumes that the opponent has a five in the hole, fall between **Tight** and **Loose.** More conservative strategies are, *ceteris paribus,* less likely to bet in particular game situations; they bluff less and require more pronounced advantages in the cards before raising or initiating betting.

In the Table 2 results, **Tight** beats **Middle** which beats **Loose.** This suggests that, *ceteris paribus,* conservative strategies may be superior performers because they risk less on each play of the game. Closer scrutiny, however, yields a qualification. Table 3 shows that **Tight's** performance falls off dramatically as the maximum bet in the game decreases relative to the minimum bet. The decrease in maximum bet lowers the amount **Tight** wins when it wins, without changing the amount of the ante (minimum bet). In effect, the ante becomes more expensive relative to prospective winnings. The superiority of conservative actions depends on their costlessness relative to prospective winnings.

Comparison of the relative conservatism (Table 1) of all nonlearning strategies with their performance (Table 2) reveals that the strategies that perform best overall— **CalcMuch, Middle** and **AvgHand**—are moderate. Conservative is not always better, even with the cost qualification. The desired degree of aggressiveness of particular actions depends on the particular context. A more skilled player should *learn* about both the game and the opponent.

3.6. Learning

Learning can occur in several ways. One type of learning adjusts a strategy's parameters. For example, a strategy may keep a conservatism index on each opponent.

When the opponent acts on a particular hand, the learning player observes the aggressiveness of the act and updates the parameter for this opponent in the appropriate direction. A second type of learning involves choosing among already known rules. For example, as a strategy learns about a particular opponent, it may change the proportions with which the elementary strategies in a complex strategy are played. A third type of learning involves the creation of new rules or the elaboration or elimination of old rules. This is the most difficult sort of learning to simulate or incorporate into models because it requires a model that somehow "understands" the game's domain and contains mechanisms for generating new strategies. We briefly examine only the first two types of learning. In the first type of learning, parameters within rules are adjusted to improve performance. **PlayerCalc,** for example, keeps historical data on its opponents and uses it to adjust rule parameters. **PlayerCalc** observes things like "whenever player #77 has bet in the past, he has had at least a ten-point advantage over what he sees in my up cards." This information is then used to refine winning probability estimates and improve performance. However, players like **PlayerCalc,** who watch the play of the opponent and try to draw conclusions about the value of the opponent's down card are vulnerable to bluffing in a way that, for example, **CalcMuch** is not.

Should a player play against players or play only against cards? There are advantages both ways. Clearly, playing against players makes **PlayerCalc** a formidable strategy against consistent players. **PlayerCalc** can spot consistency and use it to advantage. However, **PlayerCalc,** because it plays against players, can be misdirected. Table 4 restates the tournament results between **PlayerCalc** and **MixedCalc. MixedCalc** is the same strategy as **CalcMuch,** except that **MixedCalc** plays the random strategy **Simple** about 6 times in every 100. This small amount of bluffing confuses **PlayerCalc** to the point where it plays no better than **CalcMuch** against **MixedCalc.**

There is a cost to bluffing.[4] The "confused" **PlayerCalc** still beats **MixedCalc,** although not as badly as it

Table 3	Tight's Sensitivity to Maxbet/minbet Ratio							
	T	M	L	T	M	L	T	M	L
T	–	91	99	–	70	100	–	5	100
M	9	–	96	30	–	98	95	–	100
L	1	4	–	0	2	–	0	0	–
max/min		10/5			8/5			6/5	

---

[4] By "bluffing" we here mean something different than in the context of the strategy BluffsLots. Bluffing here means adding randomness to

Table 4	The Effects of Bluffing		
	PC	CM	MC
PC	–	93	61
CM	7	–	62
MC	39	38	–

beats **CalcMuch**. **MixedCalc** loses to **CalcMuch** because they are the same strategy except that **MixedCalc** plays as a much less-skilled strategy 6 percent of the time.

Game theorists have long puzzled on how much bluffing is a good (optimal) amount. Clearly a strategy that takes only random actions in order not to reveal any information to the opponent will not be a good strategy. **Simple's** poor performance is a good illustration. Just as clearly, a strategy that never bluffs will be beaten by a learning strategy capable of recognizing and exploiting patterns in play. The question of the optimal amount of bluffing has meaning only in a specific game against specific opposition.

Strategies like **CalcMuch** that play only against cards cannot be bluffed because they ignore opponents' actions. In competition where there is inept bluffing that might be discovered and exploited, ignoring opponents is costly. However, in competitions where there is effective bluffing, the level of the best play-against-cards strategy may be very close to the best that can be achieved. If there is so much noise that inference becomes infeasible given realistic constraints on player capabilities, then playing only against cards may be an excellent strategy. The overall strength of **PlayerCalcB**, the bluffing version of **PlayerCalc**, supports this contention. **PlayerCalcB** bluffs while observing the opponent's play to decide whether or not to play against cards or the opponent. If the opponent is bluffing effectively, **PlayerCalcB** reverts to playing against cards. Note that when **PlayerCalcB** is bluffed, it reverts to a bluffing ver-

sion of **CalcMuch**—it becomes **MixedCalc**. As nearly equivalent strategies, **PlayerCalcB** and **MixedCalc** play each other somewhat evenly in the tournament.

The second type of learning, choosing a strategy from among known rules, can be represented (and studied) by creating complex strategies that play each of their simple strategies with some probability adjusted on experience in play. A complex strategy, **Tight-Middle-BluffsLots (TMB),** can beat each of the three simple strategies comprising it individually. The **TMB** strategy begins by playing each of the three simple strategies with equal probability, then it adjusts the probabilities to reflect success in play. For example, if the opponent is **Middle,** then **Tight** (usually) becomes the predominant simple strategy played by **TMB**. Creating a mechanism which allows the simulated **TMB** strategy to converge to the appropriate simple strategy (e.g., **Tight** against the opponent **Middle**) is nontrivial. Convergence difficulties are relevant to real game outcomes. For example, if **TMB** plays against **Middle** and is "unlucky" in its first several instances of playing **Tight** against the opponent (if **Middle** has particularly good cards on those occasions, for instance), then **TMB** may adjust probabilities so that it rarely plays as **Tight.** Thus, it may learn to play **Tight** against **Middle** too slowly to avoid ruin.

The random device in skill-chance games greatly complicates learning about one's opponent. Learning strategies must reflect sampling, among other things, to be effective. There is a tradeoff between responsiveness and sensitivity to noise. Fast-learning strategies tend to over-react to spurious data in the form of insufficient samples while slow-learning strategies under-react. When data come from an opponent who may well be sending inaccurate signals in the presence of the random device, learning is often very difficult. Another sampling issue is, "How much history is relevant to learning when playing against a skillful opponent? A skillful opponent is likely to make adjustments as you do, so history that seemed relevant when you were playing differently may no longer be relevant. Imagine a **PlayerCalc-MixedCalc-CalcMuch (PMC)** playing against **PlayerCalc**. If the convergence mechanism works correctly, **PMC** will play predominantly as **MixedCalc** at first against a **PlayerCalc** opponent. When **PlayerCalc** is bluffed to the point where it has

---

a player's behavior to thwart other players' learning schemes. BluffsLots, in contrast, behaves consistently in a way that has been characterized as bluffing (e.g., betting aggressively when up cards are strong) by poker players. However, because BluffsLots plays consistently, thus conveying a pattern of play to opponents with learning ability, it is a "naive" bluffer.

reverted to playing as **CalcMuch,** then **PMC** will switch to playing predominantly **PlayerCalc,** which should soundly defeat the bluffed **PlayerCalc.** Interestingly, what this means is that **PlayerCalc** would be better against **PMC** if it could randomly *forget* its experiences a small percentage of the time.
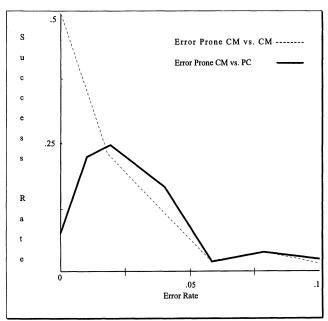
The obstacles to learning are often overwhelming in many real games with human players. In a game with a stochastic component and a task of any complexity, the number of independent variables that might be tracked quickly exceeds human capacity, even if adequate induction methods are assumed. A shrewd playing-against-cards strategy may, as we saw above in modestly complex, two-person Sum Poker, be among the best possible strategies.

3.7. Execution Skill—Reliability

Unreliable execution of rules or actions recommended by rules should impair performance. The simulation tournament confirmed this with an interesting twist. Unintentional inconsistencies in play, highly likely with human players, are as important as planned inconsistencies in making the opponent's inferences in the game environment more difficult. Figure 1 shows that introducing a small probability that **CalcMuch** will miscalculate its probability of winning actually improves its performance against **PlayerCalc.** The reason is clear upon reflection—errors are externally identical to bluffs. The scenario summarized in Figure 1 (where the vertical axis is the success rate of the error prone CM, and the horizontal axis is the rate of errors) is very similar to the scenario in Table 4 where **MixedCalc** improves against **PlayerCalc** by bluffing. In both cases, bluffing/errors with any frequency hurt the strategy against errorless and bluff-proof **CalcMuch.**

There is an optimal level of inconsistency in a strategy that depends on the learning mechanism of the opponent. Error rates that are too high begin to work against **CalcMuch** when it plays **PlayerCalc.** For example, if a 10 percent error rate is introduced into **CalcMuch,** the effects of the bluffing are overwhelmed by the force of the mistakes and error-prone **CalcMuch** becomes worse than flawless **CalcMuch** against **PlayerCalc.**[5]

---

[5] There is a growing literature on learning models in game theory (Roth and Erev 1995, Mailath et al. 1992) that discusses related issues.

Figure 1   Mistakes and Bluffing

3.8. Strategic Skill

Strategies thus far have been required to play in every game in the tournament. For many games and situations modeled as games (e.g., business competitions) this is unrealistic. What happens if strategies have the option of not playing against certain opponents?

Table 5 shows the outcome of a pairwise three-player contest in which players could choose, at any time after 10 games were played, not to play the remaining games against the opponent. For purposes of the simulation, all three strategies were given the same strategic skill criterion, namely, that they would refuse to play if their percentage of games won dipped below 10 percent.

Compare the winning percentage column (Pct.) to the net winnings column ($). Although **PlayerCalc** had the highest winning percentage, **CalcMuch** made the most money. This is because **ExpVal** is willing to play against **CalcMuch,** but not against **PlayerCalc. PlayerCalc** is, in a sense, too good, so good that it cannot find as many games in which to play. A version of **PlayerCalc** that knew that other strategies would quit if it won more than 90 percent of the time would benefit from arranging to win exactly 90-ε percent of the time. This aspect of strategic skill, inducing other players, through decep-

Table 5   Strategic Skill Tournament

	PC	CM	EV	Pct.	$'s
PC		10	10	0.95	9.5K
CM	0		77	0.7	22K
EV	1	23		0.22	−31.5K

tion or appeal to ego, to enter games where their game skills are inadequate deserves future study as "the art of the hustle."

3.9. Human Skill and Rational Strategies

None of the twelve strategies analyzed is a perfectly *rational strategy* for playing sum poker by maximizing expected value. While finding a best response strategy to a particular strategy known to be used by the opponent is conceivable, there is no reason to suppose that such a strategy is robust against mistaking the opponent's strategy. As Borel (1924, Trans. 1953, p. 115) put it: "The player who does not observe the psychology of his partner and does not modify his manner of playing must necessarily lose against an adversary whose mind is sufficiently flexible to vary his play while taking account of the adversary . . . There is no doubt that if the player follows strictly all the rules of an excellent treatise, and if his adversary knows it, that adversary can win by appropriately modifying his manner of play."

A Bayesian scheme would require a model of the opponent to generate likelihoods. The complexity of the strategies imaginable in this game makes this task quite difficult. As play proceeds and data about the opponent are gathered, various hypotheses concerning the "type" of the opponent might emerge from the data. Notice, however, that this calls on the bettor to have a means of generating as well as evaluating hypotheses on the basis of gathered data. In gathering data, the bettor might also need to have a plan for making his own decisions in the tree in such a way as to cause the opponent to make choices that convey information, Thus, the view of Kadane and Larkey (1982) is difficult to implement here.

The strategies ExpVal and ExpValB are extremely rough approximations. They do not involve an explicit model of the opponent. They use a very simple historical summary based on the smallest down card on which the opponent has ever bet and an historical estimate of the pot. Given these limitations, it is not too surprising that those strategies do not do very well.

Obviously, human players do not perform such elaborate data gathering and computation when they play poker. Nevertheless, some manage to be very good players, at least relative to their competition. It is beyond the scope here to study the play of human subjects and we will not speculate on what strategies humans might discover and how they might perform against the computer strategies.

4. Conclusion

The primary purpose of this exploratory paper is to understand the concept of skill in games as an initial step toward building theories of real games capable of both predicting outcomes and advising play. The paper explores aspects of the concept of skill in games with a moderately complex game, Sum Poker, for which strategies were created in the form of computer programs. The strategies are exercised in tournaments to examine the concepts of skill in games and skill differences among strategies in some detail. The mode of analysis has had much more in common with cognitive science and artificial intelligence than with mathematical game theory.

The two major premises of this work are: (1) skill and skill differences among players are important features of real games; and (2) theories of real games that aspire to predict outcomes or to aid in play must represent skill and skill differences among players explicitly. Both premises seem obviously true. It is difficult to think of any "real game" in which there are not important, persistent skill differences among players. It is difficult to imagine how one could predict outcomes or aid play without representing player skill. Representing player skill, even for a relatively simple game such as Sum Poker, is a messy task. Process representations of skill, describing players in terms of their method of play, require both a thorough understanding of the game and a detailed procedural description of the strategies employed. Performance representations of skill, characterizing players in terms of game outcomes in a history of play, require a thorough understanding of game contexts, random devices, and player behaviors as deter-

minants of success. When players' skills are represented, the resulting "game" is not obviously amenable to the usual modes of game theoretic analyses. How can skill concepts be incorporated in game theory?

4.1. Skill and Game Theory

The analysis of skill is a natural complement to the methodology offered by game theory. Combining these nominally disparate approaches is challenging but promises much more useful theories of real games.

An adequate representation necessitates, at minimum, modeling the possibility that the players are uncertain about each other's skill level and about the strategies the other is more likely to use. This naturally calls for the theory of games of incomplete information: each player is of a certain type which is, in this case, a strategy or a mixture of strategies. Each player knows his own type but has only some subjective beliefs on the type of his opponent—the now widely used model of incomplete information games proposed by Harsanyi (1967–1968). Recalling that the situation consists of multistage interactions, the even more appropriate setup is that of repeated games of incomplete information [see Aumann and Maschler (1995); Mertens (1987); Mertens, Sorin, and Zamir (1995); and many others]. Such a game provides a useful paradigm to model and study mutual beliefs about types and the evolution of these beliefs as the game proceeds. The updating of beliefs takes place in view of new information gathered mainly by observing the moves of the opponent. In such models one can speak of the optimal rate of releasing information about your type and the optimal rate of gathering information about your opponent's type; such information is usually costly, and the tradeoff is between immediate cost and future benefit.

Analyses of player skill such as for the game of Sum Poker may provide a manageable set of types (or strategies) for any specific game. The types can then be incorporated into game theoretic models to study their interactions in a given environment. For instance, if one can derive a relatively small set of skill and behavior types in a specific business activity, it might be very interesting and relevant to study various game theoretic models involving the interaction of those types. Such a combination of approaches may well lead to useful predictions and recommended behaviors.

4.2. Extensions

In addition to the development and use of player types in more traditional game theoretic models, there are many possible extensions of this work to more fully understand the concept of skill in games. One direction is to explore the Sum Poker domain more thoroughly. It will be interesting to have human subjects play against other human subjects and against the various computer strategies. The strategies utilized in this paper were not taken directly from playing experience, but were initially deduced and then adapted through preliminary experiments. Human subjects may find different and, perhaps, better strategies for playing Sum Poker. Also, for finding better strategies, it may be useful to run a tournament along the lines of Axlerod's work on iterative prisoner's dilemma.

We have explored only a handful of strategies in a limited number of competitive situations. Simple extensions to the existing framework include exploring 3- to n-player versions of Sum Poker and varying such factors as seating order and other tournament conditions.

There is a large amount of work ahead in understanding what skill is and how it can be represented in many different types of games. First, there is a large number of real games, including politics and business, where differential player skill is important but where skill is neither understood nor explicitly represented in models of these games. Some real games such as Chess, Go, tennis, golf, and bowling explicitly measure skill and acknowledge differences among players; in a few games such as golf there is elaborate "handicapping" to create fair contests among unequally skilled players. These games utilizing handicaps are particularly interesting for study because of the explicit performance measures of differential player skill.

Second, there is a large number of contrived games that have been thoroughly studied by game theorists without reference to explicit skill differences. There is a large amount of potentially useful theoretical work in adding skill to these more traditional, well-studied games and in understanding how the introduction of skill differences among players changes the analysis.

Appendix: Technical Descriptions of Sum Poker Players
This appendix describes the 12 players used in this study in terms of their condition-action rules. Rules are numbered in order of prece-

dence. Nine of the strategies are simple; Players 9, 10, and 12, **MixedCalc, PlayerCalcB,** and **ExpValB,** are complex strategies. In order to investigate certain specific issues, the paper uses simple combinations of the basic strategies to make more complex strategies.

The following shorthand notation will be used in describing rules:

CARDTOT = the deciding player's card total;

UPTOT = the deciding player's up card total;

DOWN = the deciding player's down card;

OPUPTOT = the opponent's up card total; and

$CARDVAL_i$ = face value of $card_i$.

Player 1: Simple
Flip a coin once at beginning of hand.
(1) If heads, then raise (or call after three raises);
(2) otherwise fold (or check if that is an option).

Player 2: AvgHand
(1) If CARDTOT > 1.2 × an average hand, then raise;
(2) If 1.2 × an average hand ≥ CARDTOT ≥ an average hand, then call;
(3) otherwise fold (or check if that is an option).

An average hand is 13.08 in the first round (two cards) and 19.62 in the second round (three cards).

Players 3, 4, & 5: Loose, Middle, and Tight[6]
Compute opponent's hand total by assuming the down card is a 1, 5, or 10 for **Loose, Middle,** and **Tight,** respectively: (1) if CARDTOT > opponent's total, then raise; (2) if CARDTOT = opponent's total, then call; (3) otherwise fold (or check if that is an option).

Player 6: BluffsLots
Compute: C = (UPTOT − OPUPTOT) + (DOWN − AvgLiveCard − 1) where

$$AvgLiveCard = \left(\sum_{deck} (CARDVAL_i \times I_{[card\ still\ in\ deck]}) \right) \Big/ \#\ cards\ still\ in\ deck.$$

(1) if UPTOT ≥ OPUPTOT then raise;

(2) if C > 0 then raise;
(3) otherwise fold (or check if that is an option).
$I_{[card\ still\ in\ deck]}$ is an indicator function equal to 0 if the card is showing in the hand, 1 if it is still in the deck.

Player 7: CalcMuch
Compute: Tiecard = CARDTOT − OPUPTOT.
Compute: # of winning cards still in deck
$= \Sigma_{deck} / {}_{[card\ still\ in\ deck,\ CARDVAL_i < Tiecard]};$
Compute: Pr(win) = # winning cards still in deck/# cards still in deck;[7]
(1) If Pr(win) > 0.75 then raise;
(2) If 0.75 ≥ prob(win) ≥ 0.5 then call;
(3) Otherwise fold (or check if that is an option).

Player 8: PlayerCalc
Recall from historical data the lowest value of down card that the opponent has ever been willing to bet on given the current differences in up cards—call this LOWDOWN. LOWDOWN is set = 11 until the opponent bets; then it takes on the historical value.

Compute: Tiecard = CARDTOT − OPUPPTOT.

Compute: # winning cards still in deck = $\Sigma_{deck} I_{[card\ still\ in\ deck, card < Tiecard, card > LOWDOWN]}.$

Compute: # Tiecards still in deck = $\Sigma_{deck} I_{[card\ still\ in\ deck, card = Tiecard, card > LOWDOWN]}.$

Compute: Pr(win) = # winning cards still in deck/# cards still in deck.

Compute: Pr(tie) = # Tiecards still in deck/# cards still in deck.

Compute: Pr(loss) = 1 − Pr(win) − Pr(tie).
(1) If prob(win) > 0.75, then raise;
(2) If 0.75 ≥ prob(win) ≥ 0.5, then call;
(3) Otherwise fold (or check if that is an option).

Player 9: MixedCalc
This player plays as **CalcMuch** with probability 0.94 and as **Simple** with probability 0.06.

Player 10: PlayerCalcB
This player plays as **PlayerCalc** with probability 0.94 and as **Simple** with probability 0.06.

---

[6] Players 4 through 10 check to see if they are "iced out" at the beginning of the second betting round (after the last card is dealt). A player is iced out when any opponent's up cards sum to more than the player's hand total. The player in question then recognizes that he cannot win and folds at the first opportunity.

[7] CalcMuch only approximates winning probability in first round, by assuming that the first round is the last round and computing winning probability accordingly. Unlike the more exact first round probability, this approximation could conceivably be calculated by a human player.

Player 11: ExpVal

Recall from historical data the average amount the opponent can be expected to contribute to pot for the remainder of the hand given the current difference in up cards—call this MATCH AMT. Compute Final Pot = 2 × MATCH AMT if neither player folds.

Compute Cost of Staying = max{MATCH AMT − amount already contributed, 0}.

Recall from historical data the lowest value of down card that the opponent has ever been willing to bet on given the current differences in up cards—call this LOWDOWN. LOWDOWN is set = 11 until the opponent bets; then it takes on the historical value.

Compute TieCard, # winning cards still in deck, # of tie cards still in deck, Pr(win), Pr(tie), Pr(loss), exactly as in **PlayerCalc** strategy.

Compute:

$$EV = [Pr(win)$$
$$\times (FinalPot + amount\ already\ contributed)]$$
$$+ [Pr(tie) \times ((0.5 \times FinalPot)$$
$$+ amount\ already\ contributed)]$$
$$- [Pr(loss) \times Cost\ of\ Staying].$$

Define AvgUp as 6.5 in first betting round, 13.0 in second.

1) If $EV > 0$ and UPTOT > AvgUp, then raise;
2) If $EV > 0$ and UPTOT ≥ AvgUp, then bet;
3) Otherwise fold (or pass if that is an option).

Player 12: ExpValB

This player plays as **ExpVal** with probability 0.94 and as **Simple** with probability 0.06.

References

Anderson, J. R., *Cognitive Psychology and its Implications*, Freeman, San Francisco, CA, 1980.

Aumann, R. J. and M. Maschler, *Repeated Games with Incomplete Information*, MIT Press, Cambridge, MA, 1995.

Axelrod, R., *The Evolution of Cooperation*, Basic Books, New York, 1984.

—— and W. D. Hamilton, "The Evolution of Cooperation," *Science*, 211, (1981), 1390–1396.

Beasley, J. D., *The Mathematics of Games*, Oxford University Press, New York, 1990.

Binmore, K., *Essays on the Foundations of Game Theory*, Basil-Blackwell, Cambridge, MA, 1990.

Borel, E., "On Games That Involve Chance and the Skill of the Players," *Econometrica*, 21 (1953), 101–115 (translated by L. J. Savage, University of Chicago, from "Sur les jeux ou interviennent l'hasard et l'habilete des joueurs," *Theorie des Probabilities*, Librairie Scientifique, Paris, 1924).

Camerer, C., "Behavioral Game Theory," in R. M. Hogarth, Ed., *Insights in Decision Making: A Tribute to Hillel J. Einhorn*, University of Chicago Press, Chicago, IL, 1990.

Caulkins, J., A. Barnett, P. D. Larkey, Y. Yuan, and J. Guranson, "The On-Time Machines: Some Analyses of Airline Punctuality," *Oper. Res.*, 41, 4, (1993), 710–720.

Gardner, M., "On the Paradoxical Situations that Arise from Nontransitive Relations," Mathematical Games, *Scientific American*, 231, 4, (1974), 120–125.

Harsanyi, J. C., "Games of Incomplete Information Played by Bayesian Players," parts I, II, III, *Management Sci.*, 14 (1967–1968).

Jacoby, O., *Oswald Jacoby on Gambling*, Hart Publishing Company, New York, 1963.

Kadane, J. B., "Does Electronic Draw Poker Require Skill to Play?" in M. H. DeGroot, S. E. Feinberg, and J. B. Kadane, Eds., *Statistics and the Law*, Wiley, New York, 1986, 333–343.

—— and P. Larkey, "Subjective Probability and the Theory of Games," *Management Sci.*, 28, 2, (1982), 113–120.

—— and T. Seidenfeld, "Equilibrium, Common Knowledge, and Optimal Sequential Decisions," in C. Bicchieri and M. L. Dalla Chiara, Eds., *Knowledge, Belief, and Strategic Interaction*, Cambridge University Press, Cambridge, 1992.

Kalai, E., "Bounded Rationality and Strategic Complexity in Repeated Games," in T. Ichiishi, A. Neyman and Y. Tauman, Eds. *Game Theory and Applications*, Academic Press, San Diego, CA, 1990, 131–157.

Larkey, P., "How to Measure Strength of Field," *Golf World*, March 29, 1991.

—— and J. Caulkins, "Incentives to Fail," H. John Heinz III School of Public Policy and Management, working paper, 1992.

Leifer, E. M., *Actors as Observers: A Theory of Skill in Social Relationships*, New York, Garland Publishing, 1991.

Livingston, A. D., *Poker Strategy and Winning Play*, J. B. Lippincott Company, Philadelphia, 1971.

Mailath, G., "Introduction: Symposium on Evolutionary Game Theory," *J. Economic Theory*, 57, 2 (1992), 259–277.

Mertens, J. F., "Repeated games," in *Proceedings of the International Congress of Mathematicians (Berkeley)*, American Mathematical Society, Providence, RI, 1987.

——, S. Sorin, and S. Zamir, "Repeated Games," CORE Discussion Papers 9420, 9421, 9422, CORE Belgium, 1995.

—— and S. Zamir, "Minmax and Maxmin of Repeated Games with Incomplete Information," *International Journal of Game Theory*, 9, (1980), 201–215.

Morehead, A. H., *The Complete Guide to Winning Poker*, Simon and Schuster, New York, 1967.

Roth, A. and I. Erev, "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8 (1995), 164–212.

Scarne, J., *Scarne's Guide to Modern Poker*, A Fireside Book, New York, 1980.

Simon, H., "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69, (1955), 99–118.

—— and W. G. Chase, "Skill in Chess," *American Scientist*, 61, (1973), 394–403.

——, "From Substantive to Procedural Rationality," in S. Latis, ed., *Method and Appraisal in Economics*, Cambridge University Press, Cambridge, UK, 1976.

——, "How Big is a Chunk?" *Science*, 183, (1974), 482–488.

——, *Models of Man*, (see especially Chapter 16, "A Comparison of Game Theory and Learning Theory"), John Wiley & Sons, Inc., New York, 1957. (*Psychometrika*, 21, (1956), 267–272.)

——, "Organizations and Markets," *Journal of Economic Perspectives*, 5, 2, (1991), 25–44.

——, *Reason in Human Affairs*, Stanford University Press, Stanford, California, 1983.

——, "Theories of Decision Making in Economics," *American Economic Review*, 49, (1959), 253–283.

Thackery, T. Jr., *Dealer's Choice*, Henry Regnery Company, Chicago, 1971.

von Neumann, J. and O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, NJ, 1944.