Tilburg University

Measuring Skill in More-Person Games with Applications to Poker

Hendrickx, R.L.P.; Borm, P.E.M.; van der Genugten, B.B.; Hilbers, P.

Publication date:
2008

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Hendrickx, R. L. P., Borm, P. E. M., van der Genugten, B. B., & Hilbers, P. (2008). *Measuring Skill in More-Person Games with Applications to Poker*. (CentER Discussion Paper; Vol. 2008-106). Microeconomics.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

No. 2008–106

**MEASURING SKILL IN MORE-PERSON GAMES WITH
APPLICATIONS TO POKER**

By Ruud Hendrickx, Peter Borm, Ben van der Genugten,
Pim Hilbers

December 2008

ISSN 0924-7815

Measuring skill in more-person games with applications to poker

Ruud Hendrickx^{1,2} Peter Borm¹ Ben van der Genugten¹
Pim Hilbers¹

Abstract

In several jurisdictions, commercially exploiting a game of chance (rather than skill) is subject to a licensing regime. It is obvious that roulette is a game of chance and chess a game of skill, but the law does not provide a precise description of where the boundary between the two classes is drawn. We build upon the framework of Borm and Van der Genugten (2001) and Dreef et al. (2004) and propose a modification of the skill concept for more-person games. We apply our new skill measure to a simplified version of poker called Straight Poker and conclude that this game should be classified as a game of skill.

Key words: games of chance, games of skill, poker

JEL classification number: C72

1 Introduction

In many countries, commercially exploiting games is subject to legal restrictions. Usually, the law makes a distinction between games of skill and games of chance. Whereas in most jurisdictions one is free to organise a chess tournament, starting a casino offering games like roulette is subject to regulation. For many games, however, it is not immediately clear to which class they belong and the way a game of chance is defined by the law only provides a partial answer.

The motivation for this research comes from the Dutch Gaming Act (1964), which states in Article 1 that

[...] it is not allowed to: exploit games with monetary prizes if the participants in general do not have a predominant influence on the winning possibilities, unless in compliance to this act, a licence is granted [...].

¹CentER and Department of Econometrics and Operations Research, Tilburg University.

²Corresponding author: PO Box 90153, 5000 LE Tilburg, The Netherlands. E-mail: ruud@uvt.nl.

To fix terminology, a game of chance is defined as a game satisfying the condition of this Act and a game of skill is any other game. In this paper, we restrict our attention to the Dutch situation, although in several other jurisdictions a similar definition of a game of chance is used and our analysis can be applied there without modification.

In a game of chance, the players do not, by definition, have a *predominant* influence on the winning probabilities. Hence the need arises for a quantitative assessment of the skill involved in playing a particular game. Borm and Van der Genugten (2001) present a model measuring the skill of a strategic game, *ie*, a game in which the outcome is determined by the players' strategy choices and not, *eg*, on their physical abilities or encyclopedic knowledge. The game's outcome however can also be influenced by external chance elements. In the basic framework the probabilities involved are assumed to be known (think of drawing cards from a deck), although using statistical techniques, the model can be extended to incorporate estimated probabilities, as is done in Van der Genugten et al. (2004).

The skill measure by Borm and Van der Genugten (2001) measures the relative skill of a strategic game on a scale from 0, which corresponds to a game of chance, to 1, which corresponds to a game of skill. The underlying idea is that in a strategic game, a player's payoff is determined on the one hand by how skillfully he plays and on the other hand by random factors. The relative weight of these effects determines the game's skill. Dreef et al. (2003) studies the relative skill measure of the simple two-player poker game with alternate bidding of Von Neumann and Morgenstern (1944).

Dreef et al. (2004) modified the skill measure introduced in Borm and Van der Genugten (2001) in order to capture the effect of internal random moves stemming from players' behavioral choices. In the current paper we propose a further modification with the aim of reducing computational complexity, which allows us to analyse more elaborate and realistic more-person games like poker. This new setting, which boils down to keeping the strategies of a player's opponents fixed throughout the analysis, has the further advantage of being more natural and transparent than the more complex behaviour of the opponents as modelled in the original setting. It is important to note however that for one-player games, the modifications have no effect on the skill level. In section 2 we present and discuss the new relative skill measure. An extensive overview of skill measures and related literature can be found in Dreef (2005).

Because the number of strategies in a typical poker game is huge, special techniques are needed to cut down on both memory usage and computing time.

In section 3 we illustrate and discuss several computational aspects of a game’s relative skill level in general. In section 4 we apply the relative skill measure to a stylised version of poker called Straight Poker. The results indicate that the relative skill level of Straight Poker lies substantially above the 0.2 threshold, which legal precedence suggests marks the boundary between games of chance and games of skill.

2 Skill in games

The Gaming Act makes a distinction between a player’s actions as a determinant of his winning possibilities (measured in terms of monetary gain) and extraneous factors. We call the effect of a player’s strategy the *learning effect* and the effect of the extraneous (chance) elements the *random effect*. To quantify these effects, we define three player types. A *beginner* is a player who has just mastered the rules of the game and is endowed with a particular (typically naive) strategy, that possibly involves randomisation (a mixed strategy). An *optimal player* is a player who completely mastered the rules of the game and picks a strategy that maximises his expected gains. A *fictive player* is a player who chooses a gain-maximising strategy, whilst knowing in advance the realisations of all random moves. We discuss the exact extent of this knowledge of random moves in more detail towards the end of this section.

For each of the three player types, the player’s strategy may depend on his position or role in the game. In a poker game, players behave differently in the first position at the table than they do at the last position. The strategy of a player typically depends on which position at the table he occupies.

In a one-player game, like roulette or blackjack¹, the expected gain of a player in each role is completely determined by the definitions above. For a beginner, you simply compute the expected gain of his endowed strategy, while for the optimal and fictive player, you solve a constrained optimisation problem to determine their optimal behaviour. In a game with more players, however, a player’s gain is ambiguous, because you should specify *against which* (strategy of the) opponents the gain should be computed.

Borm and Van der Genugten (2001) assume all opponents of a player i to jointly play a strategy which is minimax in the associated two-player zero-sum game in which they form a coalition against player i . Dreef et al. (2004) assume all opponents of player i to form a coalition and play a joint best response against

¹Both roulette and blackjack are of course played with more players, but a particular player’s gain does not depend on the other players’ strategy choices. So, from a strategic point of view, these games can be viewed as a series of parallel one-player games.

player i 's strategy. Contrary to this, we instead make the assumption that all opponents are beginners, each endowed with a predetermined strategy. We have three reasons for this. First of all, beginners form a natural benchmark. The learning effect is supposed to measure the effect of mastering all the strategic intricacies of the game and the most natural way to measure this is to keep all other things (*ie*, the opponents' strategies) constant. Second, the minimax strategy of the opposing coalition as needed in the original approach need not be unique, leading to indeterminacy of the skill level. Finally, and perhaps most importantly, with multiple players game trees are large even in the most basic variant of poker, so calculating mutual best responses becomes impossible and assuming a fixed beginner's strategy for each opponent is a necessity from a computational perspective.

We denote the finite set of roles by R and the gain of the beginner, optimal player and fictive player in role $r \in R$ by g_r^b , g_r^o and g_r^f , respectively. In our analysis of poker, we take the gain to be the player's expected payoff.

The learning effect in role $r \in R$ is defined as the difference in gain between the optimal player and the beginner

$$LE_r = g_r^o - g_r^b$$

and the random effect is defined as the difference between the gain of the fictive player and the optimal player

$$RE_r = g_r^f - g_r^o.$$

It follows from the definitions of the player types that both effects are always non-negative.

Next, we average the effects over all player roles:

$$LE = \frac{1}{|R|} \sum_{r \in R} LE_r, \quad RE = \frac{1}{|R|} \sum_{r \in R} RE_r.$$

If the learning effect is not predominant, *ie*, small compared to the random effect, the game is deemed a game of chance. This leads to the following definition of relative skill S :

$$S = \frac{LE}{LE + RE}.$$

A skill level of 0 indicates a pure game of chance in the sense that the beginner and the optimal player have the same gain. A skill level of 1 indicates a pure game of skill, because apparently the fictive player cannot obtain additional

gains compared to the optimal player using the extra information he has about the chance moves. In particular, the latter occurs in games that possess no chance elements (like chess).

Dreef et al. (2004) argue, contrary to Borm and Van der Genugten (2001), that the fictive player should be endowed with knowledge of the realisations of internal chance moves (an opponent randomising between various pure strategies) as well as external chance moves (*eg*, the dealing of cards). We elaborate on this assumption within our new framework in which all opponents are assumed to be beginners.

Consider the well-known game of stone-paper-scissors. Two players simultaneously choose either stone, paper or scissors by means of a gesture. A player choosing stone beats his opponent if he chooses scissors, scissors beats paper and paper beats stone. If both players choose the same, the game ends in a tie. Because of the cyclical winning condition, there is no *a priori* distinction between the three pure strategies. As a result, it seems reasonable to take the strategy in which stone, paper and scissors are all played with probability $\frac{1}{3}$ as the beginner's strategy. The optimal player can not play this game any better than the beginner, both having an expected gain of 0.

Because this game has no external chance moves, the fictive player as modelled in Borm and Van der Genugten (2001) also arrives at the same result. Consequently, $LE = RE = 0$ and the relative skill level is undetermined. Intuitively, however, one would say that stone-paper-scissors is a pure game of chance. The way to incorporate this intuition into the model is to fully capture *all* chance elements in the definition of the fictive player by assuming that he also knows beforehand the realisations of the opponents' internal chance moves. In stone-paper-scissors, this would lead to a positive expected gain for the fictive player (indeed, he will win every single game), leading to a positive random effect and a relative skill level of 0.

To give an impression of the magnitude of S , the table below provides an overview of the relative skill level of various one-player games (*cf* Dreef (2005)). Again note that our proposed modifications play no role in one-player games.

Game	S
Standard roulette	0
American roulette	0.004
Golden Ten	0.012
Blackjack	0.06

A recent case with far-reaching consequences involved Grand Prix Manager

2003 (GPM 2003). This is a so-called management game, in which a participant acts as the manager of a fictive sports team. The goal is to assemble a motor racing team (in terms of car components and personnel) that performs well in a simulated season of Formula One motor racing. Note that a management game is not a strategic game, since the probabilities involved in the external chance moves are not explicitly known to the players. As argued in Van der Genugten et al. (2004), however, because the game has many participants and many rounds, one can use statistical techniques to analyse the skill level of such a management game as if it were a strategic game.

Two of the current authors were expert witnesses in the GPM 2003 case. Van der Genugten et al. (2004) determined the relative skill level of various variants of GPM 2003, depending on the exact prize scheme. Using a gradual scheme (*ie*, one in which the prizes are not restricted to only the top few players in the final ranking), the relative skill level of GPM 2003 equals approximately 0.3. Comparing this with earlier verdicts on games that were judged to be games of chance, it was argued that a reasonable threshold above which a game should be considered a game of skill would be 0.2. Arnhem District Court accepted the report in full (2 February 2005, nr.105364), thereby setting a legal precedent for our skill level threshold estimate of 0.2.

3 Computing gains

As discussed in the previous section, in order to compute the skill measure of a game, you have to determine the (expected) gains of a beginner, an optimal player and a fictive player, all playing against beginners. In this section we discuss some related computational aspects.

For the beginner, the computations are straightforward. In each role, the beginner is endowed with a predetermined strategy, possibly mixed. Note that in the context of poker, a player’s pure strategy is a function that assigns an action to each information set and can hence depend on that player’s private information (in particular, the cards he holds). So for each realisation of the external chance moves, we know all the probabilities on the actions in the game tree, as well as the gain in each leaf.

For expositional purposes we do not analyse a poker game in this section, but rather a simpler fictional game which is sufficiently general and in which the calculations can be illustrated more clearly. The only external chance element in this two-player zero-sum game is that player 2 receives either hand A or hand B, each occurring with probability $\frac{1}{2}$. We illustrate the calculations for each of

the three player types in the role of player 1. There is no variation in the hand of player 1 throughout this section: for a player of any type, to compute the expected gain in a game where he can hold one of several hands, you perform the calculations for each hand he might hold separately, conditioning all payoffs and probabilities on the event that he receives this hand, and then take the average, weighted with the probabilities of each hand occurring.

In case player 1 is a beginner, he bases his strategy, by assumption of naive play, only on his own hand and not on the hand player 2 holds. In case he is a fictive player, he knows which hand player 2 holds. In both cases, to compute the expected gain, you can perform the calculations for each of the opponent's hands separately and then take the average. So for our purposes it suffices to illustrate the analysis for the beginner and fictive player only in case player 2 holds hand A. The optimal player 1 does not know which hand player 2 holds, but by inference he possesses some partial information which he should use to determine a best response. As a consequence, we have to perform the analysis for both hands simultaneously to obtain the proper conditional probabilities.

In Figure 1 we depict the game tree, where in each node there is a choice between action left (L) and right (R). The beginners' mixed strategies are indicated by the probabilities on the arcs and the italic numbers represent the gains of player 1 (equalling the losses of player 2) on the leaves. Player 2's strategy and the gains are both conditional on player 2 holding hand A. The random moves in the various nodes are assumed to be independent and the probabilities depicted in the tree are all conditional on the corresponding choice node being reached.

In Table 1 we give an overview of all the probabilities on the leaves and the corresponding contributions to the expected gain of player 1. The total expected gain of player 1 if his opponent holds hand A equals $8\frac{43}{120}$.

Leaf	5	6	8	9	12	13	15	16	20	21	23	24	27	28	30	31	total
Gain	1	3	5	7	9	11	13	15	2	4	6	8	10	12	14	16	
Prob.	$\frac{1}{48}$	$\frac{5}{48}$	$\frac{1}{12}$	$\frac{1}{24}$	$\frac{1}{10}$	$\frac{2}{5}$	0	0	$\frac{1}{40}$	$\frac{1}{40}$	$\frac{1}{80}$	$\frac{7}{80}$	$\frac{1}{40}$	$\frac{1}{40}$	$\frac{1}{60}$	$\frac{1}{30}$	1
Exp. gain	$\frac{1}{48}$	$\frac{5}{16}$	$\frac{5}{12}$	$\frac{7}{24}$	$\frac{9}{10}$	$\frac{22}{5}$	0	0	$\frac{1}{20}$	$\frac{1}{10}$	$\frac{3}{40}$	$\frac{7}{10}$	$\frac{1}{40}$	$\frac{3}{10}$	$\frac{7}{30}$	$\frac{8}{15}$	$\frac{1003}{120}$

Table 1: Probabilities of the leaves and expected gain (beginner vs beginner)

If player 1 is a fictive player, we have to compute for each realisation of the internal chance moves of player 2 (still a beginner holding hand A) a best response. In our stylised game we have to perform these calculations only twice, once for each possible hand of the opponent. In a real poker game, however,

Figure 1: Beginner (1) versus beginner (2) with hand A

the number of possible realisations of the external chance moves over which the average should subsequently be taken is typically very large. As we argue in the next section, that number can be reduced without too much loss of accuracy by considering a small number of equivalence classes (categories). For the internal chance moves, however, there is no such obvious *a priori* reduction. Instead, we speed up calculations by observing that the gain at each leaf only depends on the external chance moves. Knowing the realisations of the internal chance moves, the fictive player knows which leaves he can reach by choosing the appropriate actions. For each realisation, he chooses a strategy which gives him the highest gain. The crucial point now is that some leaves are picked for multiple realisations. The leaf with the highest gain for the fictive player is chosen whenever the realisations of the opponents' choices allow it to be reached, the leaf with the second-highest gain is chosen if it is reachable and the leaf with the highest gain is not, and so on.

So, for each realisation of the external chance moves, in our case player 2 holding hand A, we recursively determine the leaf with the highest gain for the fictive player, compute the probability of it being made reachable by the

internal chance moves of the opponents, conditional on the event that no leaf of a previous iteration is chosen, and compute the corresponding contribution to the expected gain of the fictive player. This recursion continues until all leaves that are chosen with positive probability have been dealt with.

In Figure 2 we have drawn the game tree again, where the probabilities of player 2's fixed strategy given hand A are shown on the arcs and the gains of the fictive player 1 again are shown at the leaves.

Figure 2: Fictive player (1) versus beginner (2) with hand A

By “ $(n : L)$ ” with n a decision node for player 2, we denote the event that at n , L is chosen and by “ $(n : R)$ ” the event that R is chosen. By “ $(\neq n)$ ” with n a leaf, we denote the event that player 2's choices are such that n is not reachable.

Node 31 has the highest gain for player 1 among all leaves that are reachable with a positive probability (Step 0 in Table 2). If the beginner plays R at 17 and R at 29, then node 31 can be reached. The probability that the beginner plays such that 31 can be reached then equals

$$P((17 : R) \cap (29 : R)) = P((17 : R))P((29 : R) | (17 : R)) = \frac{4}{15}.$$

If node 31 can be reached, the fictive player will always select it, regardless of what other nodes can be reached, because no other node offers a higher payoff. Thus we know that in $4/15$ of all possible realisations of the internal chance moves, the fictive player will be able to select node 31 and that, by assumption of him playing a best response given his fictive information, he will do so.

Next we determine what happens if node 31 is not reachable. This event, $(\neq 31) = ((17 : R) \cap (29 : R))^c$, happens with probability $1 - \frac{4}{15} = \frac{11}{15}$. We proceed by conditioning all probabilities in the game tree on this event. Nodes 17 and 29 are the only choices affected by this conditioning, which in turn affect the conditional probabilities of leaves 20–31. For node 29 we have

$$\begin{aligned} P((29 : R) | (\neq 31)) &= \frac{P((29 : R) \cap ((17 : R) \cap (29 : R))^c)}{P(((17 : R) \cap (29 : R))^c)} \\ &= \frac{P((29 : R) \cap (29 : R)^c)}{P((29 : R)^c)} = 0, \end{aligned}$$

where the second equality follows from $(29 : R) \subset (17 : R)$. Similarly, $P((29 : L) | (\neq 31)) = 1$.

At node 17, we have

$$\begin{aligned} P((17 : R) | (\neq 31)) &= \frac{P((17 : R) \cap ((17 : R) \cap (29 : R))^c)}{P(((17 : R) \cap (29 : R))^c)} \\ &= \frac{P((17 : R) \cap (29 : R)^c)}{P(((17 : R) \cap (29 : R))^c)} \\ &= \frac{P((17 : R))P((29 : L) | (17 : R))}{1 - P((17 : R))P((29 : R) | (17 : R))} \\ &= \frac{\frac{2}{5} \cdot \frac{1}{3}}{1 - \frac{2}{5} \cdot \frac{2}{3}} = \frac{2}{11}. \end{aligned}$$

Given the new conditional probabilities on the edges, we next compute for each leaf the probabilities of it being reachable, conditional on $(\neq 31)$. The results are in Table 2 (Step 1).

Because node 16 has the highest gain among the leaves that can be reached with positive probability conditional on $(\neq 31)$, the fictive player will select that node. The probability of this given that node 31 is not reachable is $5/12$. So the absolute probability that the fictive player selects node 16 equals $P(16 | (\neq 31)) * P((\neq 31)) = 5/12 * 11/15 = 11/36$. Next, we condition the probabilities on the event $(\neq 16)$ similarly as in the first step, only now for nodes 2 and 14 instead of 17 and 29. This yields the probabilities conditional on $(\neq 31) \cup (\neq 16)$ presented in Table 2 (Step 2).

We continue with this procedure until all the (absolute) probabilities of the

Leaf	5	6	8	9	12	13	15	16	20	21	23	24	27	28	30	31
Gain	1	3	5	7	9	11	13	15	2	4	6	8	10	12	14	16
Step 0	$\frac{1}{18}$	$\frac{5}{18}$	$\frac{2}{9}$	$\frac{1}{9}$	$\frac{2}{15}$	$\frac{8}{15}$	$\frac{1}{4}$	$\frac{5}{12}$	$\frac{3}{10}$	$\frac{3}{10}$	$\frac{3}{40}$	$\frac{21}{40}$	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{2}{15}$	$\frac{4}{15}$
Step 1	$\frac{1}{18}$	$\frac{5}{18}$	$\frac{2}{9}$	$\frac{1}{9}$	$\frac{2}{15}$	$\frac{8}{15}$	$\frac{1}{4}$	$\frac{5}{12}$	$\frac{9}{22}$	$\frac{9}{22}$	$\frac{9}{88}$	$\frac{63}{88}$	$\frac{1}{11}$	$\frac{1}{11}$	$\frac{2}{11}$	0
Step 2	$\frac{2}{21}$	$\frac{10}{21}$	$\frac{8}{21}$	$\frac{4}{21}$	$\frac{3}{35}$	$\frac{12}{35}$	$\frac{3}{7}$	0	$\frac{9}{22}$	$\frac{9}{22}$	$\frac{9}{88}$	$\frac{63}{88}$	$\frac{1}{11}$	$\frac{1}{11}$	$\frac{2}{11}$	0

Table 2: Conditional probabilities of the leaves

leaves chosen by player 1 add up to 1. In this instance, this occurs in 7 steps, which are summarised in Table 3.

Step	leaf picked	gain	prob.	exp. gain
1	31	16	$\frac{4}{15}$	$\frac{64}{15}$
2	16	15	$\frac{11}{36}$	$\frac{55}{12}$
3	30	14	$\frac{7}{90}$	$\frac{49}{45}$
4	15	13	$\frac{3}{20}$	$\frac{39}{20}$
5	24	8	$\frac{7}{40}$	$\frac{7}{5}$
6	9	7	$\frac{1}{120}$	$\frac{7}{120}$
7	23	6	$\frac{1}{60}$	$\frac{1}{10}$
sum			1	$\frac{4841}{360}$

Table 3: Overview of procedure to determine the fictive player’s gain

The expected gain for a fictive player 1 in case his opponent holds hand A is finally computed by taking the weighted average of the gains at the seven chosen leaves, yielding $13\frac{161}{360}$.

A more straightforward way of determining the fictive player’s gain would have been to enumerate all realisations of the internal chance moves and for each realisation determine the gain by backward induction. But our approach speeds things up considerably by enumerating only the fictive player’s best responses and conditioning the corresponding probabilities accordingly. Moreover, we are

only interested in his *expected* gain anyway and not in the actual corresponding strategy profile.

Next, consider the case in which player 1 is an optimal player. In order to determine his best response given his private information (cards), we apply backward induction. The external chance moves and internal chance moves determine a probability distribution on all the opponents' actions in the game tree. As a result, for each node of the optimal player that can thus be reached with positive probability, we can compute for each realisation of the external chance move the conditional probability of it having occurred and hence, the conditional distribution of payoffs at the leaves. In our stylised game we therefore cannot show the analysis only for player 2 holding hand A separately, but in order to capture the conditioning properly we have to perform the calculations for both hands of player 2 simultaneously. Starting at the final nodes of the optimal player 1, we determine the action with the highest expected payoff and from there on work our way back through the game tree.

To illustrate the computations for the optimal player 1, we can again consider Figure 2 in case player 2 holds hand A and Figure 3 in case player 2 holds hand B. Note that both player 2's strategy and the payoffs at the leaves depend on player 2's hand.

To determine the optimal player's best response, we first determine what he should do in nodes 3, 10, 18 and 25. Denote by A the event that player 2 holds hand A and by B the event that he holds B. Then, in node 3 the conditional probability that the optimal player assigns to A equals, using Bayes' law,

$$\begin{aligned} P(A | (2 : L)) &= \frac{P((2 : L) | A)P(A)}{P((2 : L) | A)P(A) + P((2 : L) | B)P(B)} \\ &= \frac{\frac{1}{3} \cdot \frac{1}{2}}{\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{3} \cdot \frac{1}{2}} = \frac{2}{5}. \end{aligned}$$

Similarly, you can compute all conditional probabilities on both hands in all four aforementioned nodes. With these probabilities, player 1 can determine in each node the expected gain of his actions. If in node 3 he plays L , his expected gain will be

$$\frac{2}{5} \cdot \frac{1}{6} \cdot 1 + \frac{2}{5} \cdot \frac{5}{6} \cdot 3 + \frac{3}{5} \cdot \frac{1}{4} \cdot 2 + \frac{3}{5} \cdot \frac{3}{4} \cdot 3 = \frac{163}{60}.$$

In Table 4 we summarise player 1's decisions in nodes 3, 10, 18 and 25.

Next, we determine what an optimal player 1 should do in node 1. Since in node 1, he does not have any additional information regarding his opponent's hand, the conditional probability of hand A in node 1 simply equals the a priori

Figure 3: Optimal player (1) versus beginner (2) with hand B

Node	3	10	18	25
Cond.prob. A	$\frac{2}{5}$	$\frac{4}{7}$	$\frac{6}{11}$	$\frac{4}{9}$
Cond.prob. B	$\frac{3}{5}$	$\frac{3}{7}$	$\frac{5}{11}$	$\frac{5}{9}$
Exp.gain L	$\frac{163}{60}$	$\frac{338}{35}$	$\frac{268}{33}$	$\frac{307}{27}$
Exp.gain R	$\frac{203}{30}$	$\frac{153}{14}$	$\frac{21}{4}$	$\frac{359}{27}$
Choice	R	R	L	R

Table 4: Optimal player's choices

probability of $\frac{1}{2}$. Using this, and taking into account the choices presented in Table 4, the expected gain if player 1 plays L in node 1 equals

$$\frac{1}{2} \left(\frac{1}{3} \cdot \frac{203}{30} + \frac{2}{3} \cdot \frac{153}{14} \right) + \frac{1}{2} \left(\frac{1}{2} \cdot \frac{203}{30} + \frac{1}{2} \cdot \frac{153}{14} \right) = \frac{331}{36}$$

and the expected gain of playing R equals

$$\frac{1}{2} \left(\frac{3}{5} \cdot \frac{268}{33} + \frac{2}{5} \cdot \frac{359}{27} \right) + \frac{1}{2} \left(\frac{1}{2} \cdot \frac{268}{33} + \frac{1}{2} \cdot \frac{359}{27} \right) = \frac{209}{20},$$

so player 1 will choose R in node 1 with an expected gain of $10\frac{9}{20}$.

4 Skill in Straight Poker

In Straight Poker, each player is dealt a 5-card hand from a standard deck of 52 after everyone has put a predetermined ante in the pot. The players then decide in turn whether to pass (at no cost) or bet (costing a predetermined bet size), until everyone has passed or one player has bet. If all players pass, everyone gets his ante back and the payoff is zero. If some player bets, then all other players (including the ones who passed before) in turn get the choice to fold, call or raise. If a player folds, he is no longer in play and loses his ante. If a player calls, he has to match the amount of money put in the pot by the previous player who did not fold. If a player raises, he matches the amount by the previous non-folded player and puts in an extra bet size. The number of raises that can be made is bounded by a predetermined maximum. Of course, once this ceiling has been reached, the remaining players can only choose between fold and call. When all players have either folded or called, the non-folded players show their hands and the player with the highest-ranking hand² wins the pot. The pot is split in case more than one player has the highest-ranking hand. In one variant, the casino always takes a fixed percentage of the pot, called the *rake*.

In order to determine the relative skill level of Straight Poker, we have to compute the strategies of the optimal player and the fictive player against a beginner. The problem is that the strategy spaces of the players in this game are huge. Depending on the exact specification of the rules, the game tree can have many nodes. Moreover, since a player's strategy is a function of the cards he holds, we should perform calculations for any possible card combination.

In order to reduce the latter source of complexity, we partition the set of poker hands into equivalence classes, called *categories*. Following, *eg*, Billings et al. (2003), it seems reasonable to assume that a player takes the same action when he holds, *eg*, $\clubsuit 9 \diamondsuit 9 7 3 2$ as when he holds $\spadesuit 9 4 \heartsuit 9 7 6$. For simplicity, we impose the same partition into categories for all three player types, although we will only use a coarsening to describe the beginner's strategy. An overview of the categories is given in Table 5, where we list the best hand in each category.

²For an overview of the ranking of poker hands, see, *eg*, http://en.wikipedia.org/wiki/Hand_rankings.

In the first 32 categories, the five cards are unsuited (*ie*, not *all* of the same suit). Any hand stronger than two pairs (including a flush) is contained in category 33.

In order to describe a reasonable beginner’s strategy for Straight Poker, we start with two assumptions. First, a beginner will make a decision on whether to pass or bet or whether to fold, call or raise only on the basis of the cards he holds and not on the role he occupies. Second, a beginner will not use a bluffing strategy: he will bet with stronger hands and pass with weaker hands. Likewise, he will raise with strong hands, call with intermediate hands and fold with weak hands. So, to fully describe a beginner’s strategy, we only have to specify three boundaries.

In our analysis, we consider two different beginner’s strategies, which we present in Table 6. For each boundary we indicate the highest category in which the “lower” action is chosen. So, in both strategies a beginner bets with a pair or better and raises with a pair of jacks or better. In strategy 1, the beginner folds whenever he has no pair and in strategy 2, he folds whenever he has less than ace-jack high.

Note that both beginner’s strategies are pure. In a subsequent sensitivity analysis, however, we also consider mixed strategies.

To get an impression about which variables are important for the relative skill level, we first analyse two-player Straight Poker, before proceeding to the more computationally intensive variants with more players. The ante is set to 1 and there is no rake. For both cases of the beginner’s strategy we computed the relative skill level for bet sizes 2, 4 and 8 and a maximum number of allowed raises of 1, 2 and 3. In practical poker variants, a bet size of 2 or 4 and a maximum of 3 raises is common. The results, which are taken from Hilbers (2007), are presented in Table 7.

From the results in Table 7 we conclude that even in the most simple variant of Straight Poker, the relative skill level lies above the critical threshold value of 0.2 discussed in section 2. The relative skill level is increasing in the bet size, whereas the maximum number of raises and the particular beginner’s strategy against which all gains are computed do not seem to be very influential.

The positive influence of the bet size on the relative skill level can be understood as follows. Compare a game with bet size 2 to a game with bet size 8, all other things being equal. The strategy of the fictive player will be the same in both games, since it only depends on the ordering of the leaves in terms of gains, not on the actual numbers involved. This implies that each time the beginner places a bet, the difference in the fictive player’s (non-negative) gain is

Cat.	card 1	card 2	card 3	card 4	card 5	hands
1	9	8	7	6	4	53,040
2	10	9	8	7	5	70,380
3	J	9	8	7	6	71,400
4	J	10	9	8	6	56,100
5	Q	9	8	7	6	71,400
6	Q	10	9	8	7	57,120
7	Q	J	10	9	7	84,660
8	K	9	8	7	6	71,400
9	K	10	9	8	7	57,120
10	K	J	10	9	8	85,680
11	K	Q	J	10	8	121,380
12	A	9	8	7	6	70,380
13	A	10	9	8	7	57,120
14	A	J	10	9	8	85,680
15	A	Q	J	10	9	122,400
16	A	K	10	9	8	85,680
17	A	K	Q	J	9	81,600
18	2	2	A	K	Q	84,480
19	3	3	A	K	Q	84,480
20	4	4	A	K	Q	84,480
21	5	5	A	K	Q	84,480
22	6	6	A	K	Q	84,480
23	7	7	A	K	Q	84,480
24	8	8	A	K	Q	84,480
25	9	9	A	K	Q	84,480
26	10	10	A	K	Q	84,480
27	J	J	A	K	Q	84,480
28	Q	Q	A	K	J	84,480
29	K	K	A	Q	J	84,480
30	A	A	K	Q	J	84,480
31	10	10	9	9	A	57,024
32	A	A	K	K	Q	66,528
33	all	hands	better	than	2 pairs	74,628
total						2,598,960

Table 5: Categories of poker hands (best hand in each category displayed)

a factor 4. The optimal player on the other hand bases his strategy on *expected* gain and different bet sizes may lead to different best responses. For hands with

Strategy	pass/bet	fold/call	call/raise
1	17	17	26
2	17	13	26

Table 6: Beginner’s strategies in Straight Poker

Beg. strat.	bet size	allowed raises	rel. skill	beg. strat.	bet size	allowed raises	rel. skill
1	2	1	0.35	2	2	1	0.34
1	2	2	0.35	2	2	2	0.35
1	2	3	0.36	2	2	3	0.35
1	4	1	0.40	2	4	1	0.40
1	4	2	0.39	2	4	2	0.40
1	4	3	0.40	2	4	3	0.41
1	8	1	0.45	2	8	1	0.46
1	8	2	0.43	2	8	2	0.45
1	8	3	0.44	2	8	3	0.46

Table 7: Skill of 2-player Straight Poker

a positive expected gain, he is confronted with the same difference in expected gain as the fictive player, but for hands with a negative expected gain, his losses will be equal. So, on average the optimal player will be better off.

To measure the sensitivity of the relative skill level with respect to the beginner’s strategy, we consider some slight alterations of the boundaries listed in Table 6. For the call/raise boundary of 26 (a pair of jacks) we performed the calculations for the alternatives 25 (a pair of tens) and 27 (a pair of queens) and the two intermediate cases in which with a pair of tens (or jacks, respectively) the beginner calls and raises with probability $\frac{1}{2}$ (and, of course, raises in each higher category and calls or folds in each lower category). For the boundary of 13 (ace-jack high) we computed the seven variants in between a 50-50 choice in category 12 (ace-nine high) and a 50-50 choice in category 15 (ace-queen high). For the boundary of 17 (a pair of twos), we computed the seven alternatives in between a 50-50 choice with 16 (ace-king high) and a 50-50 choice with 19 (a pair of threes).

We computed the relative skill level for each combination of boundaries, bet sizes and maximum number of raises. The histograms of the results for beginner’s strategies 1 and 2 are presented in Figure 4 and 5, respectively.

We observe that the variation in skill level is quite large. In all variants,

Figure 4: Relative skill histogram measured against variations of beginner's strategy 1

Figure 5: Relative skill histogram measured against variations of beginner's strategy 2

however, the skill level lies above the 0.2 threshold.

Above, we argued that the bet size is an important determining factor for

the relative skill level. If you take a closer look at the results, this finding is confirmed when one considers the variations in the beginner's strategy. Indeed, if you take only one particular bet size, the variation in relative skill level is far less pronounced. In Figure 6 we depict the histogram for the variants of strategy 1 and bet size equal to 4 (and 1, 2 or 3 raises allowed).

Figure 6: Relative skill histogram measured against variations of beginner's strategy 1, with bet size 4

The results for bet size 2 and bet size 8 are similar to Figure 6, with similar spread, and peaks at approximately 0.33 and 0.42, respectively. For each bet size, the low observations correspond to a beginner's strategy with a relatively high fold/call boundary. This confirms the earlier finding that relative skill is lower when measured against a more conservative beginner's strategy.

If we add a rake of 5% of the pot, there are two effects. The main effect is that the fictive player is far more exposed to the rake, which in effect is paid by the winner, than the optimal player. This would suggest an increase in skill if a rake is introduced. A second effect is that the expected gain of the beginner changes. His expected payoff will decrease, but the effect may be asymmetric in the player roles, so the net effect on the relative skill level is unclear. Using a subsample of all variants of the beginner's strategy discussed before, we observe that in nearly all instances, the net effect of a 5% rake on relative skill will be positive (although in some variants, it is slightly negative), with an increase of

up to 0.05.

From the above we conclude that the relative skill of 2-player Straight Poker hovers between 0.25 and 0.45, depending on the specific rules (a higher bet size leading to a more skillful game). The relative skill level depends on the beginner’s strategy against which it is measured, but in no variant does it drop below the 0.2 threshold. In almost all variants, a rake of 5% leads to an increase in relative skill.

When we consider Straight Poker with more than two players, both memory usage and computing time increase quickly. Both the tree size and the number of category combinations that have to be evaluated increase exponentially in the number of players. Since the tree size is determined by the rules of the game, we can only speed things up by reducing the number of categories. We consider the coarsening of the original 33 categories in Table 5 into the 9 categories indicated by the horizontal lines. Note that this particular coarsening allows for the beginner’s strategies mentioned in Table 6 to be expressed in terms of the new categories.

Since we are interested foremost in whether the relative skill level lies above the 0.2 threshold, we further save time by only performing the calculations for those worst-case variants which led to the lowest 10% in skill level in the 2-player case. So we implicitly take the view that the number of players has a negligible effect on which values of the other parameters lead to a minimum relative skill level. The results for three and four players are presented in Figures 7 and 8, respectively. These results provide an indication that with more players, the relative skill level of Straight Poker will be higher.

The conclusion is that Straight Poker, which is a rather stylised variant of fixed-limit poker, is a game of skill. Note that in no-limit poker or poker in tournament form, betting decisions are typically of a different kind, requiring a separate but similar analysis.

In poker variants that are more realistic than Straight Poker, like Texas Hold’em and 7-Card Stud, there are usually two additional sources of complexity: typically, there are many betting rounds and each player’s final (five-card) hand is composed in a more elaborate way. In a game with more moves there is more scope for the optimal player to obtain information by inference which the fictive player gets for free. So in a more complicated game the information gap between the optimal and fictive player closes and the random effect has a relatively smaller impact. One would therefore expect that the more complex the game tree becomes, the higher the relative skill level will be.

Figure 7: Relative skill, 3 players

Figure 8: Relative skill, 4 players

References

- Billings, D., N. Burch, A. Davidson, R. Holte, J. Schaeffer, T. Schauenberg, and D. Szafron (2003). Approximating game-theoretic optimal strategies for full-scale poker. In: *Proceedings of IJCAI-03*.
- Borm, P. and B. van der Genugten (2001). On a relative measure of skill for games with chance elements. *TOP*, **9**, 91–114.
- Dreef, M. (2005). *Skill and strategy in games*. Ph. D. thesis, Tilburg University, Tilburg, The Netherlands.
- Dreef, M., P. Borm, and B. van der Genugten (2003). On strategy and relative skill in poker. *International Game Theory Review*, **5**, 83–103.
- Dreef, M., P. Borm, and B. van der Genugten (2004). A new relative skill measure for games with chance elements. *Managerial and Decision Economics*, **25**, 255–264.
- Genugten, B. van der, P. Borm, and M. Dreef (2004). De toepassing van de wet op de kansspelen op de managementspelen Competitie Manager en Grand Prix Manager. Report, Department of Econometrics and OR, Tilburg University, Tilburg, The Netherlands. (In Dutch).
- Hilbers, P. (2007). Skill in poker. MPhil thesis, Department of Econometrics and Operations Research, Tilburg University, Tilburg.
- Neumann, J. von and O. Morgenstern (1944). *Theory of games and economic behavior*. Princeton: Princeton University Press.