

Journal of Expertise
2020, Vol. 3(2)
© 2020. The authors
license this article
under the terms of the
Creative Commons
Attribution 3.0 License.
ISSN 2573-2773

Poker as a Domain of Expertise

Jussi Palomäki¹, Michael Laakasuo¹, Benjamin Ultan Cowley^{1,2}, and Otto Lappi¹

¹Cognitive Science, Department of Digital Humanities, University of Helsinki, Finland

² Faculty of Educational Sciences, University of Helsinki, Finland

Correspondence: Jussi Palomaki, jussi.palomaki@helsinki.fi

Abstract

Poker is a game of skill and chance involving economic decision-making under uncertainty. It is also a complex but well-defined real-world environment with a clear rule-structure. As such, poker has strong potential as a model system for studying high-stakes, high-risk expert performance. Poker has been increasingly used as a tool to study decision-making and learning, as well as emotion self-regulation. In this review, we discuss how these studies have begun to inform us about the interaction between emotions and technical skill, and how expertise develops and depends on these two factors. Expertise in poker critically requires both mastery of the technical aspects of the game, and proficiency in emotion regulation; poker thus offers a good environment for studying these skills in controlled experimental settings of high external validity. We conclude by suggesting ideas for future research on expertise, with new insights provided by poker.

Keywords

Economic decisions, probabilistic decision-making, risk, expertise, poker

Introduction

In everyday and expert settings, humans are able to cope with high levels of complexity and ambiguity. We are able to make economic decisions under time pressure, on the basis of limited information, and with various levels of risk and uncertainty associated with the outcomes. Most of the decisions are menial, such as which type of bread to buy for dinner; others are personally and professionally significant, such as whether to trade a stock at a given price. Some decisions may even be life changing, such as deciding to undergo surgery on short notice. How humans make such decisions is a foundational issue in behavioral economics, and in social and cognitive psychology. This issue is also important for research on expertise, because some decisions (such as trading stocks) are made in a manner

that may be conducive to the development of expertise (involving repeated performance, explicit criteria for decision quality, competitive environment, and feedback).

Ultimately, to understand expertise in risky decision-making we need to discover what psychological mechanisms underpin both the success and failure of decisions in complex, ambiguous, and intricate real-world settings (Klein, 2008; 2015). Unfortunately, the settings of such real-world problems are generally not readily amenable to traditional experimental methods. Therefore, the cognitive underpinnings of human decisions are often investigated in highly simplified laboratory tasks, which are intended to capture some hypothetical mechanism or essential aspect of real-world problems (Buelow & Blaine, 2015; Buelow & Suhr, 2009; Kahneman & Tversky, 1979). This

creates a tension: Restricted tasks abstract away much of real-world domain complexity, ambiguity, and the “world knowledge” that experts¹ bring to bear on the task. This allows one to arbitrate more definitively among competing mechanistic hypotheses, but also raises the question of whether those putative mechanisms are a factor in more realistic settings.

Laboratory tasks are meant to be analogues of real-world environments, but whether the tasks actually have relevance outside the laboratory has to be taken on faith (that is, researchers’ intuition for how similar their simplified, abstract decision-making task really is to a real-world task setting). Traditional decision-making tasks are thus designed for laboratory convenience—often presented in text or numerical form using novice subjects and/or with domain-general problems. This makes them particularly limited in terms of shedding light on *skilled* decision-making processes in rich and more natural contexts.

The study of *games* has been a valuable route for cognitive scientists and can offer some middle ground between experimental control and ecological validity. Most everyday natural decisions—such as choosing ingredients for cooking a meal or deciding on what to wear to a party—cannot be given comprehensive mathematical definitions, nor are there often clear normative criteria on the goodness of a decision. However, many games are everyday tasks with definable rules that can be compactly represented. Also, gameplay offers means to design recurring situations that can be used to present decision-making tasks that have both experimental control and high ecological validity (such as choosing the next move in chess). Game decisions can, moreover, often be varied in terms of task difficulty and complexity to suit particular participants or experimental questions. Finally, mathematical analysis of games has in many cases provided normative standards whereby decision quality is evaluated.

In this review we show how these desirable characteristics apply to the game of poker, which can serve as a valuable model system for studying expert economic decision-making

under risk and uncertainty. Poker is a well-structured game played in a social setting with many different game variants involving randomness and probabilistic economic decision-making. These aspects of poker make it attractive for various scientific disciplines interested in economic or rational decision-making at the individual and social levels. Poker also comes with a very large online community of players generating big datasets and powerful data-gathering opportunities (e.g., Eil & Lien, 2014; Siler, 2010) similar to many electronic sports (Esports) games (e.g., Thompson, McColeman, Stepanova, & Blair, 2017). In general, games that have gone online provide enormous research opportunities—poker in particular, given its long history of defined analytic structure, game theoretical analysis, as well as large player base.

We will argue that poker also offers a novel look into expertise, since the concept of poker skill is more complex than the much-studied *technical* skill in other well-studied game domains such as chess. This is due to the element of chance in the game: Skilled poker players need to have *emotional* tolerance of *outcome variability*—that is, to be successful they need to be able to control and reflect on their negative emotions when even right choices can lead to catastrophically bad outcomes merely due to chance. Compared with chess, poker is also typically played with a larger group of people, making emotion regulation particularly important.

So far, this element of poker has not been thoroughly studied, despite the potentially significant benefits for decision- and cognitive sciences. Therefore, poker has strong, but as yet untapped, potential for research on social- and cognitive psychology, decision-making, and expert performance.

Overall, while poker has received a lot of attention outside academia², up until recently much of the research on poker has been *clinically* motivated; for example, evaluating how pathological gambling behavior manifests in poker, or *theoretically* focused on using poker as a testbed for artificial intelligence (see Brown & Sandholm, 2019; Moreau, Chabrol, &

Chauchard, 2016; and Rubin & Watson, 2011). We argue that games of economic decisions such as poker can and should be used more in basic behavioral research on decision-making and expertise under risk and uncertainty. The element of randomness may make mastering poker different compared to mastering many deterministic games—however, as articulated by Siler (2011), it is precisely this stochastic nature of poker that makes it a much more realistic task environment reflecting the vagaries and uncertainties of many real-life phenomena such as financial decisions.

In our review we first provide a section describing the technical aspects of poker and explain the basic structures of the game and how the element of chance influences skill development. Then we address the following review research questions (RQs): (1) *What are the components of poker skill?* We describe how the concept of poker skill comprises both

technical (mathematical, statistical and game theoretical) and *emotion regulation* (“mental game”) components, and how various social elements of the game can bias players’ decision-making. (2) *How do poker skills develop into expertise, and how does poker allow study of expertise?* We link the components and development of poker skill to previous work on *expertise, deliberate practice, and skilled intuition* and show that poker offers a novel way to look at expertise and expert performance due to its emotion regulatory skill aspects. (3) *How can future studies on expertise and decision-making make use of poker?* We conclude our review by detailing how future studies can draw insights from poker to examine skilled decision-making under emotional and social constraints.

Table 1 illustrates the features of poker reviewed in this paper and their relevance to research on expertise and decision-making.

Table 1: Features of poker and their relevance for decision-making and expertise researchers

Poker Feature	Research Relevance; Poker allows study of the following:
Incomplete information	Microcosm of <i>naturalistic</i> financial decision-making
Interplay of skill and chance	Skill perception and biases in decision-making: in the short term, bad players may win (inflated skill perception), and good players may lose (obfuscation of true skill)
Male-dominated social environment	Masculinity and gender stereotypes in a competitive setting, gender biased decision-making
Technical and emotional aspects of skill	Interplay between emotion regulation ability and decision-making accuracy
Multiple sources of both public and private information	“Game theory optimization” strategies, how skilled players avoid exploitation
Skilled intuition	Ecologically valid skilled intuition in a “medium validity” (as opposed to “high validity”; e.g., chess) environment

Basic Properties of Poker

In every poker variant the winnings of one player are the losses of another (poker is a *zero-sum* game; Wright, 2001). Decisions in poker are economic decisions made in partially unpredictable environments with potentially undesirable outcomes (it is a game of *randomness* and *risk*). Players have to decide between various options and act without seeing the other players’ cards (it is a game of

incomplete information (Sklansky & Malmuth, 1999). Players must also adapt to changes in the nature of game information across the phases of the game, and, according to Salen and Zimmerman (2004, p. 149) poker contains several *types of information*. Furthermore, while the game is turn based, the pace of game play between human opponents still often creates substantial *time pressure*³. The time used for deliberation can also indirectly disclose

information on one's strategy, further pressuring players to control their behavior. Poker is also a *dynamic* environment, as the "game state" changes even when the agent does nothing. The social pressure and the monetary stakes involved create additional cognitive and emotional load—for professional players the rewards can reach millions of US dollars.

In the technical examples that follow, we will focus on the most popular variant of poker called *No Limit Texas Hold'Em* (NLHE). In NLHE, each player is first dealt two cards, and the goal is to form the best five-card combination from one's own two cards (not seen by the other players) plus cards dealt on the table (shared with all other players). There are up to four rounds of betting, during which the number of publicly shared cards increases, starting with no shared cards and ending with, at most, five. Between each round the players can make investment decisions on whether to keep playing, how much to invest in the pot, or give up (i.e., fold)⁴. The pot will go to the winner (or split between winners in case of ties), who is the player with the best card combination, or the only one not to have folded.

Skill and Chance

Generally, poker is viewed as a game of both skill and chance, but the extent to which one or the other dominates is debated (Croson, Fishman, & Pope, 2008; Dedonno & Detterman, 2008; Fiedler & Rock, 2009; Levitt & Miles, 2014; Meyer, von Meduna, & Brosowski, 2013). Anecdotal evidence supports the view of poker as a game where one's skills can constantly be improved (Brunson, 2005; Sklansky & Malmuth, 1999; Tendler, 2011). The consensus view in academic discussion is that although chance plays a role in *short-term* results, with enough skill poker can be played profitably *in the long run*. Empirical support for this view comes from an analysis of 456 million online poker hands (van Loon, van den Assem, & van Dolder, 2015). Van Loon et al. (2015) created a simulation based on these data, comparing the best players with the worst ones, and found that skill starts to dominate chance when performance is assessed over about 1,500 or more hands of play (see Fiedler & Rock, 2009, for

similar results). Skill has a demonstrably significant role also in real-world poker success. Professional players are consistently more successful than amateurs at the World Series of Poker (Croson et al., 2008; Levitt & Miles, 2014).

One way to illustrate the role of chance in poker is through simulations of *outcome variability*. Players' levels of skill are reflected in their *win rate*, which is the average amount of profit over some number of played hands (usually 100; van Loon et al., 2015). The standard deviation of a player's win rate (a measure of outcome variability) can be 20 times higher than the win rate itself (Billingham et al., 2013). To illustrate, we will compare two equally skilled hypothetical players playing 200,000 hands each. By assuming both players have somewhat low win rates (on the statistical edge of making long-term profit), we might observe the situation presented in Figure 1: One player could be winning substantially (> 15 000 €), and the other clearly losing (-5000 €). Outcome variability is thus a highly significant factor, masking a player's "true" skill as defined by the expected long-run winnings (dashed line in Figure 1). This means that while poker differs from games of pure chance (such as roulette) or games of skill and chance where long-term profit is unattainable (e.g., blackjack; Bjerg, 2010), outcome variability still makes it challenging to empirically estimate the actual skill level of any individual player from naturalistic play data⁵.

However, player skill can also be estimated experimentally, by using *representative decision-making tasks*, with known normative solutions: more (technically) skilled players should consistently reach that solution more quickly and/or reliably. In two laboratory studies (Linnet et al., 2010; 2012), those who had played poker at least once a week for at least a year were better at estimating betting outcomes than less experienced ones. Two online studies with simplified poker tasks showed that the amount of poker experience was strongly and positively associated with making mathematically appropriate poker decisions (Laakasuo, Palomäki, & Salmela, 2015; Palomäki, Laakasuo, & Salmela, 2013a). Thus, components of poker skill can be isolated and studied both "in the wild" and in the laboratory.

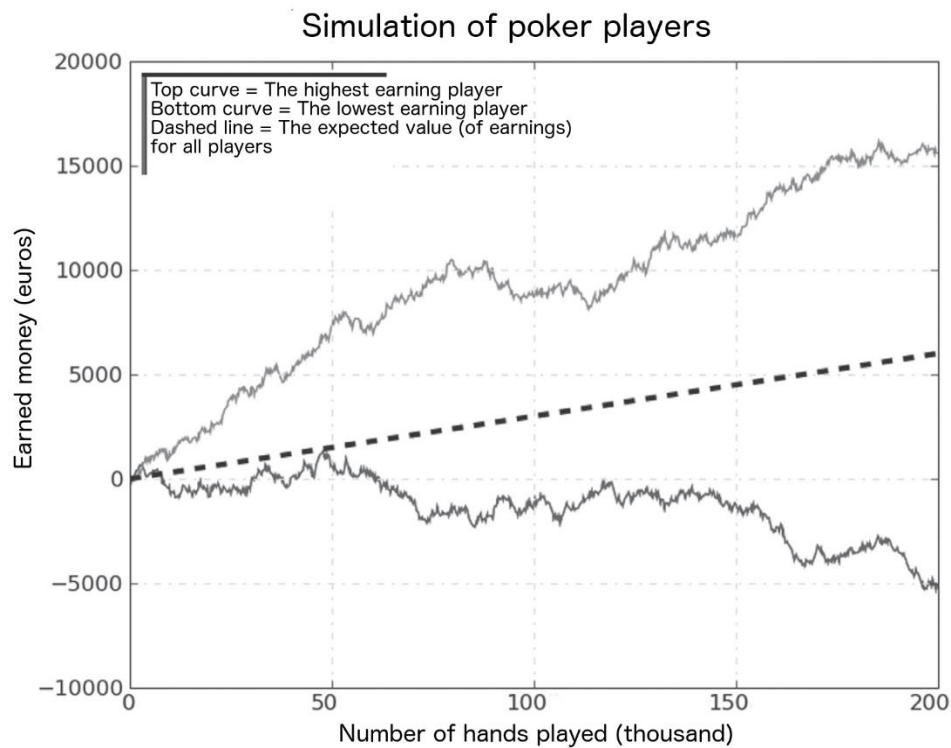


Figure 1. A simulation of 900 NLHE “poker players” with equal win rates. Win rates are based on “big blinds”; that is, the minimum bet size allowed by the rules. These win rates are calculated based on 3 big blinds—in this case, euros—per 100 hands played, with a standard deviation (*SD*) of 80 for the win rates (typical win rate *SD*s in NLHE are 70-90; Billingham et al., 2013). Note that the “players” are simulated processes based on two parameters (win rate and $SD_{win\ rate}$) and thus independent of one another. The figure depicts only the highest and lowest earning players (top and bottom curve, respectively), and the expected value of earnings for all 900 players (dashed line). (Translated into English from Laakasuo, Palomäki, & Lappi, 2015).

Components of Poker Skill (RQ1)

In this section we address our first research question on the components of poker skill. We consider what is required of a good poker player; that is, someone who is generally able to make a long-term profit by playing poker. We propose a division of poker skill into *technical* and *emotion regulatory* (sub)skills.

Technical poker skills refer to in-depth knowledge of game mechanics and betting strategies, and how to apply them to increase one’s chances of winning. In poker, technical skills alone are not enough for long-term success if dysfunctional emotional responses systematically impair players’ decision-making. Ample evidence shows that emotions have a significant impact on success in poker, and emotion regulation skills are necessary to play poker consistently at a high level. Below, we

explain how acquiring mastery of poker involves not only technical and strategic knowledge of the game but also an aspect of “mind management” or mental game ability.

Technical Skills

In terms of technical skill elements, Billings and colleagues (2002) have proposed that in order to play poker, one needs to understand at least the following concepts: (1) *hand strength and hand potential*, (2) *betting strategy, bluffing, unpredictability*, and (3) *opponent modeling*. Palomäki et al. (2013a), among others, have suggested that (4) *bankroll management* is also vitally important. These four key elements are explained below.

Hand strength and hand potential refer to how strong a player’s hand currently is and the probability of a given hand strength changing—

relative to the opponents' assumed hand strengths—as further cards are dealt (see the Appendix for detailed examples). Calculations of hand strength and hand potential require knowledge of poker betting odds in a given situation, mathematical aptitude, and working memory capacity (e.g., DeDonno, 2016; Meinz et al., 2012).

Betting strategy, bluffing, and unpredictability refer to knowledge of when and how much to bet or raise (or fold) in a sufficiently unpredictable manner to maximize one's profit and protect oneself from exploitation. Betting strategy refers for instance to the decision to bluff with a fixed frequency or not, as a player might decide *a priori* to bluff a given number of times in a game. These skill elements require players to apply (either explicitly or implicitly) the concepts of game theory, such as *Nash equilibrium*⁶, in their own decision-making.

Opponent modeling refers to estimating the full range of an opponent's possible hands. Specifically, opponent modeling relates to how various behavioral and social opponent characteristics, such as betting patterns, physical "tells", or gender, influence the way (or what range of possible hands) one's opponents are predicted to play—and, consequently, how they should be played against to maximize profit. This generally requires interpreting concealed social signals, reading covert facial expressions, and detecting deception in general.

Bankroll management is the knowledge of how much money is needed for playing, in relation to the stakes played, to avoid going broke. That is, how much capital is needed to withstand outcome variance and avoid "going broke due to merely bad luck." Good bankroll management skills are typically associated with a good understanding of the concepts of statistical variance and risk of ruin (Browne, 1989; Palomäki et al., 2013a).

The depth of the technical aspects of poker is evidenced by clear differences in technical skill between proficient and novice poker players. For example, in a laboratory experiment St. Germain and Tenenbaum (2011) compared the performance of proficient players, with significant tournament success, to intermediate

and novice poker players in a simulated poker task during which participants had to "think out loud." Proficient players outperformed both intermediate and novice players (in terms of profit), and self-reported the highest amount of thought processing and attention to relevant technical aspects of the task—such as betting patterns, estimated opponent ability, future opponent actions and "tells," and hand selection and strength. Practicing these skill elements leads to better performance: DeDonno and Detterman (2008) conducted a laboratory experiment where naïve poker players systematically practiced technical poker concepts which lead to improved success in the game. The players were given information and feedback about (1) when and why to pay attention to the other players' decisions; (2) the concept of playing fewer hands, and how to play them; and (3) hand ranking strategy with quality values for the initial hand. These correspond to opponent modeling, betting strategy, and evaluating hand strength and hand potential, respectively.

Poker is a knowledge-rich domain, with complex demands on both technical and strategic skills. These demands present information processing challenges requiring the player to go beyond the information embodied in the cards and explicit in the rules. We have provided a detailed poker task analysis in the Appendix, which illustrates the complexities involved in poker decision-making.

Poker is also well-suited to facilitate study of players' information processing because the relationship between different forms of information is relatively clear and understood. The above aspects of technical poker skill embody different challenges of *information manipulation* (Salen & Zimmerman, 2004, p. 148), in the sense of the information embodied by the cards as defined by the rules of poker. This information can take multiple forms, including: (1) *Information known to all players*; i.e., the five "community cards" shown on the table; (2) *Information known to only one player*; i.e. the two "hole" cards of each player; (3) *Information known to the game only*; e.g., the unused cards in the deck; and (4) *Randomly*

generated information; i.e., the shuffling of the deck. Part of the technical skill in poker is to know which form of information the game embodies at any given time—and to keep track of and predict how the game’s information moves between these forms.

For example, a certain amount of the information known only to one player (the hole cards), can be leaked to other players due to that player’s response to the *turn*, *flop*, or *river* (see the Appendix for explanation of these terms). In fact, in a recent study, Frey, Albino, and Williams (2018) analyzed 1.75 million poker hands and found that winning (skilled) poker players were better than losing (unskilled) players at *integrative* information processing—creating new information based on the interaction between their own hole cards and their opponents’ betting patterns. This made the winning players’ decision-making less exploitable and harder for others to reverse engineer (Frey et al., 2018).

To recap, technical poker skills consist of knowledge of hand strength and hand potential, betting and bluffing strategy, ability to avoid exploitation (playing unpredictably) and to exploit others, and bankroll management; all of which can be viewed in terms of challenges for information manipulation. However, we note that empirical research on technical poker skill development in terms of information manipulation strategies is still relatively scarce.

Emotion Regulation Skills

Due to statistical variance in the game, even technically skilled poker players regularly encounter losing streaks and “bad beats” — losing money in situations where losing is objectively unlikely, and not the result of normatively poor decision-making. Losing large sums of money often elicits negative emotions, which, in turn, can have detrimental effects on upcoming decisions. For example, in a bout of anger an experienced and otherwise technically skilled player might forgo sound betting or bankroll-management strategies, ending up playing with too high stakes and betting erratically despite *factually* knowing it is mathematically inadvisable. Thus, in addition to

technical skill elements, the concept of poker skill encompasses an *emotion regulatory* aspect. Emotion regulation skills refer to the ability to withstand the arduous, yet inevitable, losing streaks without having them affect the quality of one’s decisions (Boujou et al., 2013; Palomäki et al., 2013a). These skills may be conscious processes explicitly controlling one’s impulses by willpower or positive self-talk, or they could be more unconscious processes, which might be termed trait emotional stability or “character” developed by surviving previous encounters.

McCormack and Griffiths (2012) interviewed professional and recreational poker players and found that professional players were not only more likely to be logical and controlled in their behavior, but also took fewer risks and were less likely to chase after losses (i.e., keep playing in an attempt to win back their losses). Conversely, recreational players showed more signs of losing control, taking unnecessary risks and playing under the influence of intoxicants. In correlational online studies, poker experience has been found to be negatively associated with the psychological traits of *emotionality* (Laakasuo, Palomäki, & Salmela, 2014), *self-rumination* (Laakasuo, Palomäki, & Salmela, 2016; Palomäki et al., 2013a) and *emotional sensitivity to losses* (Laakasuo et al., 2016; Palomäki et al., 2014). Experienced players are thus less emotional, dwell less on negative thoughts, and report higher emotional tolerance of poker losses than do inexperienced players.

Moreover, Palomäki and colleagues (2013a) report that in an online setting with simplified poker tasks, experienced players—but not inexperienced—made mathematically better poker decisions when they had a strong tendency for *self-reflection*. Self-reflection is a trait related to analyzing one’s past mistakes in a cool and detached manner. Consistent with these results, Leonard and Williams (2015) employed a measure of technical poker skills and betting strategy and found that proficient players were less susceptible to gambling fallacies and had higher emotional tolerance for financial risk and better social information processing skills.

Tilt: Intense Moral Anger Revealed in Poker

The importance of emotion regulation skills and aversion to loss (via pursuing risk) in poker is underscored by the phenomenon known as *tilting*, which refers to losing control due to negative emotions—typically related to bad beats or prolonged losing streaks—and making strategically weak decisions and losing significantly more money than otherwise (Palomäki et al., 2014; Moreau, Delieuvin, Chabrol, & Chauchard, 2017). Extreme cases of tilting have even led to losing entire life savings within minutes, and to self-reported memory losses of the preceding events (Palomäki et al., 2013b; Tendler, 2011).

Poker communities seem to agree that tilting is a significant phenomenon: in an online study, 88% of poker players reported having tilted severely at least once within their last 6 months, 43% more than five times, and 24% more than 10 times (Palomäki et al., 2014). Hence, this form of “mental disarray” occurs with a substantial frequency, leading to substantial costs for those involved. These findings are in line with the studies by Smith, Levere, and Kurtzman (2009), as well as by Eil and Lien (2014), who used big data on millions of played online poker hands and found that players tend to pursue risk when losing, but play cautiously when winning. This effect is possibly driven by emotional aversion to loss.

Social cues may also interact with emotional reactions during poker decision-making: In an online experiment employing a poker decision-making task with mathematically defined optimal choices, inducing the feeling of anger (via reading emotional stories) reduced decision-making accuracy. However, this effect was driven by a social cue: displaying a pair of human eyes that “followed” the participants’ mouse cursor during the task (Laakasuo, Palomäki, & Salmela, 2015).

What leads to such costly lapses in judgment? In a qualitative study on poker players’ experiences of losing significant amounts of money, tilting was characterized by feelings of anger, frustration, and significantly, injustice (Palomäki et al., 2013b; see also Barrault et al., 2014). Social elements such as unfriendly comments by other players often fuel

the negative emotional states leading to tilting (Browne, 1989). The sense of *injustice* is particularly interesting, as it makes tilting a form of moral emotion: Individuals (sampled in Palomäki et al., 2013b) who tilt reported feeling personally insulted, and that they “unfairly” lost money for which that had worked diligently. They viewed variance as “bad luck,” took it personally, and started pouring their earnings into the game and chasing their “fair chance.” The authors postulated (Palomäki et al., 2013b) that the psychology of tilting could be viewed as *moral anger*: Losing due to bad luck is perceived as “cosmically” unjust, which motivates an overly aggressive yet ineffective retaliation strategy of excessive betting. In the aftermath of tilting, the players reported being disappointed in themselves and that they were ruminating over lost resources.

Experienced players, however, differ from inexperienced ones in their reporting of better skills for regulating negative game-induced emotions. Some experienced players in (Palomäki et al., 2013b) reported that their general emotion regulation skills had improved through playing poker. These players also thought that a clear understanding of mathematical concepts, such as variance, is related to a mature disposition towards encountering “bad luck” (“*luck doesn’t exist, only variance does*” [Palomäki et al., 2013b]), which suggests that in poker, emotion regulation skills and technical skills are tightly intertwined.

Social Nature of the Game

In poker, the dynamics of social interaction—such as opponent characteristics or gender effects—are crucial in understanding decision-making quality. The social setting of the game also plays a significant role in biasing poker decisions on the one hand, and on the other provides players with potentially accurate information in the form of behavioral “tells” (Caro, 2003).

Schlicht, Shimojo, Camerer, Battaglia, & Nakayama (2010) employed a simplified poker task and found that opponents whose facial expressions displayed more trust were more often folded (given up) against. The authors

argued that by betting the opponent is implicitly “sending a message” that he has a strong hand, and, because he looks trustworthy, the message is believed. In another study with a poker task involving repeated decisions against the same opponent, participants’ decisions were more strongly influenced by their opponents’ prior actions when the opponents were represented as humans rather than as computers (Carter et al., 2012).

In both of these studies, the human opponents were presented as males. Poker players seeking long-term engagement with the game value masculine identities and player traits (Wolkomir, 2012), and the vast majority (90–95%) of poker players are male (Palomäki et al., 2014; see also Abarbanel & Bernhard, 2012). Also, poker decision-making itself seems to be gender-biased: In an experiment using realistic online poker tasks where opponents were represented as either male or female avatars, participants (of whom 93% were male) bluffed 6% more frequently at online tables with female-only avatars compared with male-only tables amounting to a significant difference over time (Palomäki et al., 2016). A majority of the participants also reported that the gender of their opponents did not influence their decisions to bluff, which suggests an implicit (unconscious) bias in bluffing female opponents, who might have been perceived as “easier” targets than males.

Together, these results highlight the notion that the social nature of poker is a key element in fully understanding decision-making quality and biases in the game. But turning it around, poker is a tool to study decision-making and socially driven decisions in a market environment-like scenario, which, to date, has received relatively little attention in research.

Measuring Poker Skill

Time, speed, or distance measures can be used in many sports for objective quantification of performance; and in chess—the game most studied in cognitive science—Elo points provide a high-validity measure of performance. In poker, however, skill-level is often assessed indirectly by self-reported experience or

simplified poker tasks, as previously discussed. The element of chance obfuscates empirical assessment based on earnings and calls for very long observational histories. It would be better if the probabilistic “goodness” of individual players’ decisions (e.g., the expected value in terms of monetary winnings) could be evaluated based on a reasonable number of hands played.

The expected value of poker decisions can be evaluated in simplified scenarios (see Laakasuo, Palomäki, & Salmela, 2015; Leonard & Williams, 2015). However, evaluating the expected value of complex poker decisions “in the wild” is extremely difficult, given all the aforementioned cues potentially affecting (or biasing) the players’ decisions and the element of chance. One way to tackle this problem is by benchmarking poker players’ decisions against the best artificial intelligence (AI) poker programs. Somewhat recently, an NLHE AI not only won the 2016 Annual Computer Poker Competition, but in 2017 defeated four highly skilled professional poker players in heads-up (one versus one) matches for about \$1.8 million over 120 thousand hands⁷. Poker AI has thus been benchmarked against the highest human standard and proven sophisticated enough to beat the very highest-performing human players. Therefore, these programs can act as a normative reference whereby the performance of sub-elite players at least can be evaluated. This would be based on how well their decisions correspond to the consensual decisions of the best AI. To our knowledge, such efforts have not been made yet, highlighting a potential avenue for future research.

Development of Poker Expertise (RQ2)

Our second research question asked how poker skills develop into expertise and how poker allows for studying expertise. The complexity of requisite technical knowledge in poker (as explained in “Technical Skills,” above) is evident even in a simplified poker decision-making task, which we have provided in the Appendix. Poker also lends itself well to be examined under theories of *expertise*. Due to having a chance component embedded in a well-defined rule structure, poker even helps

extend existing work on expertise to domains where decision quality is not fully correlated with observed outcomes (unlike in chess, for example, where consistently making the best decisions reliably leads to good outcomes).

Deliberate Practice

The *deliberate practice* (DP) framework is the most established explanation for how expertise is acquired (Ericsson, 2007; Kaufman & Duckworth, 2017). It can be applied to study the development and acquisition of poker skill, expertise, and skilled intuition. In turn, the special features of poker relating to chance, emotion regulation, and social factors show that acquiring mastery of only the technical aspects of the game does not guarantee long-term success. So far, the DP framework has been used mainly in relation to what we have called technical skill, and therefore we suggest that the question of emotion regulatory skill development (through DP or otherwise) is an important new direction.

Within the DP framework, research on the cognitive foundations of expertise has shown that the superior performance of experts is not based on general intelligence, but on a vast amount of well-organized topic related knowledge (Ericsson, Krampe, & Tesch-Römer, 1993; Ericsson & Lehmann, 1996; Ericsson & Williams, 2007; Kaufman & Duckworth, 2017). This knowledge is clearly acquired through experience, and the DP framework characterizes the nature of that experience by making one core assumption: An individual's level of performance in the domain is monotonically related to the amount of a specific type of practice (DP) that person has engaged in. Put differently, the attained level of expertise and performance are a function of the time invested in DP. In music training DP refers to (typically solitary) practice to improve specific technical or artistic aspects of one's skill, but not studying music theory, public performances, or "jamming" (Ericsson, Krampe, & Tesch-Römer, 1993). In chess, studying and determining the best moves in mid-game⁸ would count as DP, while merely playing or spending time on studying the literature generally would not. According to

Ericsson (2016), as a predictor of performance, accumulated DP is more important than the amount of overall domain experience, general intelligence, or innate domain specific talent *combined* (for further discussion, see Ackerman, 2014; Macnamara, Hambrick, & Oswald, 2014; Macnamara, Moreau, & Hambrick, 2016; and Hambrick et al., 2014).

Technical Poker Skill Acquisition via Deliberate Practice

Although poker does not have a formal teaching culture like in classical music and professional sports, the range of self-coaching strategies suggests that the online poker sub-culture is a mature culture of expertise. A common recommendation for "serious" novice players seeking to improve their skills is to use poker analysis software, which allows for monitoring of session-by-session statistics on profit or loss and betting strategy (Billingham et al., 2013). After each session, the players can then carefully analyze how they played and what they could have done differently. Poker players actively interact over virtual communities to scrutinize poker strategy. Skilled players, in particular, frequently post detailed breakdowns of how they played for general discussion (O'Leary & Carroll, 2013). Their aim is to fine-tune their mathematically informed strategic decisions in poker (Parke & Griffiths, 2011).

We posit that in poker, this type of *study of betting strategies* in specific game situations would count as DP for technical skills (we are not aware of specific practice forms that would be geared toward emotion regularity skills, that is, DP for non-technical skills, in poker). Although this is not solitary practice designated by a teacher, the explicit *goal* of improving specific skills and the setting-up of clear *feedback* mean the process can be viewed as DP, in the context of poker.

Moreover, posting one's poker hands (breakdown of a string of decisions within a specific hand) for analysis and scrutiny on online poker forums has three characteristics of DP. First, a clear task structure, wherein the players often receive step-by-step walk-throughs on why certain decisions should or should not

be made. Such walk-throughs may also *isolate subtasks*, such as focusing on different stages of the hand (e.g., decisions on the flop, turn, or river). Second, there should be clear goals for the players who *aim for self-improvement*. For example, the feedback generated by playing the game might be positive for bad decisions (winning money despite making a decision with negative expected value), but posting such situations online for scrutiny helps players discover the actual goodness of their decisions. The proximal goal for players who seek feedback on their decisions is often not merely to enjoy poker or winning since they also post hands where they have *won* but are uncertain whether they played correctly⁹. Finally, there is the element of diligent *repetition*, as players who strive to get better keep posting poker hands for scrutiny, which, in turn, helps them increase their skills.

It should also be noted that because poker is a competitive game, skilled players might have an incentive *not* to help novice players to improve—or even an incentive to hinder their progress. Novice players aspiring to get better thus sometimes need to discern between misleading and accurate information disseminated in online poker forums (Talberg, 2019), as aspect of social skill learning.

Emotion Regulation Skill Acquisition

The consensus is that technical poker skills can be learned via practicing and studying the game. However, studying poker alone is probably not enough to learn and improve one's emotion regulation skills, because it is not easy to “simulate in training” the loss of significant amounts of money.

Traits such as low emotionality and low tendency to self-ruminate are largely (possibly innate) predispositions that enable some people to become good players; namely, those who can endure the stressful learning period as well as the unavoidable losing streaks. Personality, IQ, and other psychological traits, when measured with standard psychometric instruments, are to a large extent stable across time, and may be difficult to alter systematically through practice. However, the malleability of such traits, and the

directions of causality between poker skill development and various psychological characteristics could be fully evaluated only by employing a longitudinal study design, where poker players' behavior is measured over extended periods of time. To our knowledge, no such study currently exists and would thus offer a fruitful line for future research.

There is, however, a rich corpus of poker self-coaching textbooks that focus on improving one's mental game skills. The authors of these books typically recognize emotion control as a highly significant element in poker skill development (e.g., Angelo, 2007; Taylor & Hilger, 2007; Tendler, 2011). Similar anecdotal evidence has emerged from Esports, where professional teams and individual players have been significantly more successful in tournaments after hiring sports psychologists specializing in tilt-management (theScore esports, 2019).

Tendler (2011) draws from his experience as a clinical psychologist working extensively with poker players and offers detailed guidelines and techniques for players to improve their tilt control. He views poor tilt control in poker as an issue of consistency in individual performance level. Players perform better on some days than on others—and the overall distribution of performance level forms a bell curve around the average performance level for each player. For players with poor tilt control, this distribution is wide, reflecting a large difference in performance level between their best and worst possible performance. Players with proficient tilt control, in turn, have narrower performance level distributions. In other words, their performance is more *constant across time*—they play almost as well on their “worst day” as they do on their “best day.”

It is important to note that we do not claim emotion regulation is an important “sub-skill” *only* in the game of poker. It probably has wide relevance across a range of domains, especially those dealing with risk and uncertainty. We are, however, proposing that the role of emotion regulation becomes more pronounced in poker than most domains that have been used in cognitive science to study the nature and development of *expertise*. In other fields such as

chess or playing a musical instrument, proficiency in emotion regulation (or staying cool under pressure) might be what separates the “super-elites” from the “merely experts.” However, it is almost unheard of that poor emotion regulation skills would cause a chess grandmaster to lose against a beginner, or a professional musician to fail to perform above the novice level. In poker, however, tilting can

cause an otherwise technically proficient player to perform *extremely poorly*. This is illustrated in Figure 2, where emotion regulation skills are conceptualized to affect *within-individual variability in performance over time*—or, in other words, performance level *consistency*—for individuals of putatively identical level of technical skill.




Figure 2. Hypothetical role of emotion regulation (ER) skills in Music, Chess, and Poker among players with high *technical skills* (“on their best day”) in their respective games. The lines depict the hypothetical *within individual variation across time* in level of performance (i.e., consistency in level of performance) for individuals with (1) low ER skills in Music, Chess, and Poker (dashed lines), and (2) high ER skills in any field. For example, a technically proficient Poker player with low ER skills might sometimes perform as well as those with high ER skills, but due to high variability in their performance level, they sometimes perform as poorly as an amateur. This is not the case for technically proficient individuals with low ER skills in Chess or Music, where performance level almost never drops significantly low. Note that for simplicity, we assume that for good ER skills the profile of performance variability is the same across all fields. We also note that this is a *conceptual model*; the level of individual performance with respect to technical and ER skills in real life is likely somewhat more complex and nuanced.

Thus, some poker players who have acquired a high level of technical skills (e.g., through years of DP) might still struggle with having highly inconsistent performance levels (dashed line for Poker in Figure 2). For these players, technical skills alone are not enough to reach a high average performance level. The extent to which emotion regulation skills can be learned, and if DP would be a suitable framework to understand learning them, is

unclear and an important venue for future research.

Skilled Intuition as the Interplay of Technical and Emotional Skill

A corpus of anecdotal evidence suggests that since the poker environment is complex and fast paced, players need to trust their intuitions or “gut feelings” when making a decision (e.g., Brunson, 2005; Tendler, 2011). These feelings

can also be called affective heuristics (Finucane et al., 2000)—that is, “unconscious” processing of task-relevant information experienced phenomenologically as good or bad “feelings” about a situation. It has been empirically established that chess masters, too, often rely on an intuitive “feel” for different moves and assessment of the “board as a whole,” especially in the mid-game where options for various moves are astronomical (e.g., Chassy & Gobet, 2011; Gobet & Chassy, 2009) and it is futile to attempt to work through the alternatives step-by-step in working memory.

Cognitive modeling work suggests that chess masters’ intuition relies on pattern recognition (“chunking”). Their intuition is a cognitive process involving implicit memory and fast and automatic procedural knowledge: the present board configuration is compared to a vast knowledge-base of mid-game positions encountered over uncounted chess matches, analysis of chess literature, and thousands of hours of playing chess and solving chess problems (Gobet & Chassy, 2009). This implicit information processing seems to also be *accurate* enough to assist in complex decision-making in familiar domains. What would this type of “skilled intuition,” or more specifically implicit domain memory retrieval that benefits task performance, look like in poker? We illustrate this with a quote from a two-time World Series of Poker main event champion, Doyle Brunson (Brunson, 2005, p. 542):

Whenever I . . . “feel” . . . I recall something that happened previously. Even though I might not consciously do it, I can often recall if this same play . . . came up in the past, and what the player did or what somebody else did. So, many times I get a feeling that he’s bluffing or that I can make a play and get the pot. [My] subconscious mind is reasoning it all out.

In this quote Brunson clearly alludes to what can be called skilled intuition in the domain of poker, manifesting in episodic memory recall or gut feelings. The “feel” is, presumably, a subconscious recollection of something that has happened in the past, which cannot be

articulated in more detail. In cognitive terms, skilled intuition can be viewed as a hallmark of expert decision-makers across many domains, but it can reliably exist only in environments with *stable* relationships or regularities between identifiable cues and specific events, actions and outcomes, such as chess (Kahneman & Klein, 2009). These kinds of regular environments and games are known as *high-validity* environments. The opposite are low-validity environments, such as changes in political climates, where predictability of long-term outcomes from past performance is limited, and any intuition-based judgment is likely to be flawed or biased (Kahneman & Klein, 2009). The higher the validity of the environment, the better the chances are for acquiring skilled intuition in that environment (e.g., chess, bike riding, or the game of djenga). DP may be seen as a means to increase the validity of (some aspects of) the environment.

Is poker a high-validity environment? Given the strong element of randomness, specific decisions consistently lead to specific outcomes only over the long run. Learning poker strategy is therefore challenging because the process is masked by outcome variability (Figure 1). In the short run, players will often receive positive feedback from bad decisions, which may elicit an illusion of skill (Bjerg, 2010), and *vice versa*, obfuscating actual skill. Even after many hours of practice and play, players might have an erroneous conception of their true skill, and the soundness of their choices. Indeed, MacKay and colleagues (2014) found that increased frequency and duration of poker play was more strongly associated with online poker players’ *perceived* skill than with their *objectively measured* skill. There are tools to measure one’s level of skill objectively in chess (Elo-ratings) but not in poker; due to this poker players are also more biased in predicting their individual success in tournaments (Park & Santos-Pinto, 2010). On the other hand, poker is based on a deterministic system of rules, such that it is predictable at some scale. Thus, poker—or any other similar game where the goodness of decisions is defined only *over the long run*—

might be considered a *medium-validity* environment.

To our knowledge, no empirical studies have evaluated differences between medium- and high-validity environments in the development of expertise and skilled intuition. It is also unknown to what extent and what aspects of emotion regulation skills—which are probably more idiosyncratic to medium-validity environments—become “intuitive” (when “emotional maturity” in the face of losses is achieved), and to what extent they require constant cognitive control. During prolonged series of losses, some players—even self-proclaimed poker professionals—may start thinking the game is “rigged” against them. In other words, technically skilled players may start believing that the online poker sites deliberately manipulate who gets to win and who does not (Tendler, 2011; Palomäki et al., 2013b), even when these beliefs are not supported by evidence. Such experiences are likely more frequent in medium-validity environments than in high-validity environments (and most frequent in low-validity environments, where all kinds of irrational beliefs and “superstitions” may develop).

Poker as a research tool also offers us the chance to contrast with existing results, for example that chess masters employ different evaluation strategies to novices; they mentally falsify their hypotheses rather than confirm them as novices do (Cowley, 2017). Such comparative work would shed light on the processes by which reduction of validity affects expert decision-making strategies also for experts.

Poker as a Research Tool (RQ3)

Let us recap where we are, theoretically, before we proceed to consider specific ways poker could be used as a research tool in the study of expertise. We started by analyzing the probabilistic aspects of typical poker decisions and described the information structure of the game. We then progressed to show how this task environment is modulated by several different factors and addressed the issues related to poker skill conceptualization. Next, we placed poker within the framework of expertise and deliberate practice and suggested that

performance in poker could be largely driven by skilled intuition: Technical poker skills should not be construed just as the ability to perform explicit mental calculations, but also as the ability to make skilled intuitive judgments based on a “feel” for the game—as is also the case in more established expert domains. However, skilled intuition or gut feelings in poker may be hard to obtain due to natural outcome variability in the game (mathematically good decisions—that is, decisions with positive expected values—might not result in preferred outcomes). Also, since the poker decisions become intuitive, they are in danger of being interfered with by external factors, such as emotions of social anger and “tilting,” as well as gender stereotypes. Therefore, a crucial aspect of becoming good at poker is developing skills for regulating one’s emotions in the face of stochastic outcomes, in a challenging social environment. Based on this framework of understanding of poker, we are now in a position to illustrate some of the potential that poker holds for research on decision-making and expertise.

Generally, poker seems to be better posed for longitudinal studies than many other ecologically valid games, or purely game theoretical lab-games, since in poker the concept of skill has an important and well-defined meaning *mainly* over the long run and in the context of emotional tolerance of variance (Palomäki et al., 2013ab; Palomäki et al., 2014; Laakasuo et al., 2014). For example, the amount of DP in poker may not strongly reflect players’ long-term success unless they also invest in mental game training, which, in turn, may or may not be achievable via DP (see Figure 2). Future studies should thus look into how effective DP is the context of developing emotion regulation skills, or “mental toughness” across various fields (e.g., Tendler, 2011).

Another route for future studies is evaluating motivational factors in developing poker skills. Some research suggests that a masculine identity is very important for poker players who seek long-term engagement with the game (Wolkomir, 2012). However, we have little knowledge of how different identity factors contribute to possible

biases or errors in economic decision-making. Research on this topic is *prima facie* interesting and relevant to understanding of, for example, stock market trading or risky decision-making in general. If poker for some players is about pursuing “male glory” (Palomäki et al., 2016; Wolkomir, 2012), is it also about “manly risk taking”? Do poker players with stronger masculine identities make riskier decisions, and when are riskier decisions *better* decisions? Since poker involves risk, it would also be useful to see how risk-taking, masculine identity and emotional volatility react with and possibly hinder rational decision-making. Whether or not people make risky decisions in economics and sports for reasons of fame and glory is interesting for several reasons. Do we want egoistically-motivated stock traders, or leaders who take risks to boost their own self-image, if these motivations make them blind to disastrous outcomes?

Poker research seems to have uncovered a previously unstudied emotional state called “tilt”—a specific type of *moral* anger—which could possibly also be found in other areas of decision-making and expertise, such as sports or finance. What is going on in other domains of action when people lose control; for example, in stock market trading, golf, tennis, or racing (Wei et al., 2016)? The term “tilt” has also been adopted into common use in the world of online gaming and Esports (theScore esports, 2019). Is tilting a uniform phenomenon across of these domains; if so, how much of expertise within these domains depends on emotion regulation skills?

There seems to be a zen-like quality in top poker players who report not getting anxious

about the “swings” of their fortune in poker (Palomäki et al. 2013b), similar to the skill of experienced investors like Warren Buffet. Can we find in other domains similar results regarding emotion-regulation skill: Namely that self-reflection, emotional stability, and understanding “variance in life” (a *que sera, sera*-type stoic mentality) protects against destructive emotions? Players could be taught emotion regulation skills via, for example, mindfulness meditation (which has shown promise in improving emotion regulation in a gambling context [de Lisle et al., 2012]); meditation-based intervention might improve poker players’ decision-making.

Pinning down and measuring the elements comprising poker skills would also be informative in the study of skilled intuitions and their acquisition. The role of skilled intuition, or gut feelings, in poker decision-making has not been empirically investigated. At what level of skill do gut feelings start being accurate enough to aid in decision-making—or in other words, *when* will poker players start profiting from listening to their intuition instead of losing because of it? In poker, it is difficult to accurately estimate one’s own skill, since the observed outcomes of playing are masked by variance. This creates extra pressure on players to deliberately self-reflect on their session-by-session decisions without focusing too much on the actual results. These questions offer a fertile and significant area for study that will serve to further integrate research on emotions and decision-making. Table 2 presents our conceptualization of poker skill and its sub-components.

Table 2. Conceptualization of poker skill

Poker skill			
Technical		Emotion regulation	
<ul style="list-style-type: none"> • Understanding <ul style="list-style-type: none"> • probabilistic dependencies, chance, and variance • hand strength and hand potential concepts • bankroll management, betting strategy • opponent behavior 		<ul style="list-style-type: none"> • Tolerance for losses and “bad beats” • Responding to “swings” • Avoiding loss of control and “tilting” 	
<i>Analytic</i> Explicit step-by-step calculations in working memory	<i>Intuitive</i> Implicit assessment; affective heuristics (“gut feelings”)	<i>Impulse control</i> Cognitive control, inhibition of impulsive responses	<i>Emotional stability</i> Development of trait emotional stability, a “mature” emotional disposition

Conclusion

Poker offers an ecologically valid, rule-based, and well-structured environment of decision-making under risk and uncertainty, where decisions are made under emotional pressure in a social setting, and on the basis of substantial domain knowledge and skill. Thus, studying poker not only sheds light on human decision processes, but also on how skill and expertise in these processes develop with experience, and how social and emotional factors moderate such decisions.

Most traditional tasks used to gain knowledge on human decision-making are, in contrast, relatively simple, numerically presented, and administered in a laboratory setting. They have debatable ecological validity and may not accurately model how humans behave in more complex, naturalistic real-world settings. Moreover, the participants are almost always *inexperienced* in these tasks, and thus it is also very difficult to model how expertise would moderate any observed effects. The expertise literature, on the other hand, studies tasks and skills that are measured across domains in the real world and thus have high ecological validity, but also complexity, making them difficult to operationalize, or to determine the *a priori* normative decisions. Taking advantage of “naturally occurring” laboratories, such as poker, stands to greatly benefit the study of decision-making and expertise.

Endnotes

1. We define expertise as *the ability to reliably and consistently produce a level of performance, in a specific domain, that is much superior to the level attained by the novice.*
2. Even if poker has not been extensively studied in the area of decision-making, historically it has been important: the game is said to have inspired John von Neumann to invent game theory (von Neumann & Morgenstern, 1944).
3. In online poker, there is typically a time limit of 1-3 minutes per decision. In live poker, time constraints are not as obvious as in online poker. Nonetheless, taking “too long” to make a decision is considered bad table etiquette

(Malmuth, 2012). Moreover, in live poker, players are allowed to “call the clock” on any other player (at any time), at which point the said player typically has 60 seconds to act until his/her hand is declared “dead” (i.e. automatically folded). These rules depend on the casino where poker is played.

4. For more details, see the Appendix; for the general rules of poker consult Krieger and Bykofsky (2006).
5. More extreme win rates do not change this picture: Variance is not affected by the “degree” of win rate.
6. Nash Equilibrium in poker is when two players are playing a strategically “optimal” game (in terms of expected value) against one another, and neither can gain anything by unilaterally deviating from the said “optimal” strategy (Bowling et al., 2015).
7. See <http://www.computerpokercompetition.org> and <https://www.theguardian.com/technology/2017/jan/30/libratus-poker-artificial-intelligence-professional-human-players-competition>
8. Typically, middle game in chess is considered to begin when both players have completed the development of all or most of their pieces and the king has been brought to relative safety.
9. Note that this does *not* include situations where the player has won with an inferior hand after the odds are explicitly known – that is, due to “good luck.” Rather, here we refer to situations where, for example, the player decides to bluff and the opponent folds. In this case bluffing might actually have been incorrect, if the probability of the opponent folding was too low.

Acknowledgments

The authors thank Jane and Aatos Erkkö Foundation (grant no. 6-5291-3 to ML [PI] and JP), Emil Aaltonen Foundation, the Finnish Cultural Foundation (grant no. 00150514 to OL), and the Academy of Finland (grant no. 325694 to JP) for financial support. They also thank Drs. Jami Pekkanen, Jukka Sundvall, Eeva Palomäki and Nils Köbis for feedback on the manuscript.

Author's Declarations

The authors declare that there are no personal or financial conflicts of interest regarding the research in this article.

The authors declare that they conducted the research reported in this article in accordance with the [Ethical Principles](#) of the Journal of Expertise.

References

- Abarbanel, B. L., & Bernhard, B. J. (2012). Chicks with decks: the female lived experience in poker. *International Gambling Studies*, *12*(3), 367-385.
- Ackerman, P. (2014). Nonsense, common sense, and science of expert performance: Talent and individual differences. *Intelligence*, *45*, 6-17.
- Angelo, T. (2007). *Elements of poker*. Self published.
- Barrault, S., Untas, A., & Varescon, I. (2014). Special features of poker. *International Gambling Studies*, *14*(3), 492-504.
- Billingham, J., Cinca, E., & Tiroch, T. (2013). *The education of a modern poker player*. D&B Poker.
- Billings, D., Davidson, A., Schaeffer, J., & Szafron, D. (2002). The challenge of poker. *Artificial Intelligence*, *134*(1-2), 201-240.
- Bjerg, O. (2010). Problem gambling in poker: Money, rationality and control in a skill-based social game. *International Gambling Studies*, *10*(3), 239-254.
- Bouju, G., Grall-Bronnec, M., Quistrebert-Davanne, V., Hardouin, J. B., & Vénisse, J. L. (2013). Texas hold'em poker: a qualitative analysis of gamblers' perceptions. *Journal of Gambling Issues*, *28*, 1-28.
- Bowling, M., Burch, N., Johanson, M., & Tammelin, O. (2015). Heads-up limit hold'em poker is solved. *Science*, *347*(6218), 145-149.
- Browne, B. R. (1989). Going on tilt: Frequent poker players and control. *Journal of Gambling Behavior*, *5*(1), 3-21.
- Brown, N., & Sandholm, T. (2019). Superhuman AI for multiplayer poker. *Science*, *365*(6456), 885-890.
- Brunson, D. (2005). *Super System 2: A course in power poker*. Las Vegas, NV: Cardoza Publishing.
- Buelow, M. T., & Suhr, J. A. (2009). Construct validity of the Iowa gambling task. *Neuropsychology review*, *19*(1), 102-114.
- Buelow, M. T., & Blaine, A. L. (2015). The assessment of risky decision making: A factor analysis of performance on the Iowa Gambling Task, Balloon Analogue Risk Task, and Columbia Card Task. *Psychological assessment*, *27*(3), 777-785.
- Carter, R. M., Bowling, D. L., Reeck, C., & Huettel, S. A. (2012). A distinct role of the temporal-parietal junction in predicting socially guided decisions. *Science*, *337*(6090), 109-111.
- Caro, M. (2003). *Caro's book of poker tells*. Las Vegas, NV: Cardoza Publishing.
- Chassy, P., & Gobet, F. (2011). A hypothesis about the biological basis of expert intuition. *Review of General Psychology*, *15*(3), 198-212.
- Cowley, M. B. (2017). *Hypothesis Testing: How We Foresee Falsification in Competitive Games*. Lambert Academic Publishing 2017. Available at SSRN: <https://ssrn.com/abstract=23390>
- Croson, R., Fishman, P., & Pope, D. G. (2008). Poker superstars: Skill or luck? Similarities between golf—thought to be a game of skill—and poker. *Chance*, *21*(4), 25-28.
- DeDonno, M. A., & Detterman, D. K. (2008). Poker is a skill. *Gaming Law Review*, *12*(1), 31-36.
- DeDonno, M. A. (2016). The influence of IQ on pure discovery and guided discovery learning of a complex real-world task. *Learning and Individual Differences*, *49*, 11-16.
- Eil, D., & Lien, J. W. (2014). Staying ahead and getting even: Risk attitudes of experienced poker players. *Games and Economic Behavior*, *87*, 50-69.

- Ericsson, K. A. (2007). Deliberate practice and the modifiability of body and mind: toward a science of the structure and acquisition of expert and elite performance. *International Journal of Sport Psychology*, 38, 4-34.
- Ericsson, K. A. (2016). Summing Up Hours of Any Type of Practice Versus Identifying Optimal Practice Activities: Commentary on Macnamara, Moreau, & Hambrick (2016). *Perspectives on Psychological Science*, 11(3), 351-354.
- Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100(3), 363-406.
- Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual review of psychology*, 47(1), 273-305.
- Ericsson, K. A., & Williams, A. M. (2007). Capturing naturally occurring superior performance in the laboratory: Translational research on expert performance. *Journal of Experimental Psychology: Applied*, 13(3), 115-123.
- Fiedler, I. C., & Rock, J. P. (2009). Quantifying skill in games—Theory and empirical evidence for poker. *Gaming Law Review and Economics*, 13(1), 50-57.
- Frey, S., Albino, D. K., & Williams, P. L. (2018). Synergistic information processing encrypts strategic reasoning in poker. *Cognitive Science*, 42(5), 1457-1476.
- St. Germain, J., & Tenenbaum, G. (2011). Decision-making and thought processes among poker players. *High Ability Studies*, 22(1), 3-17.
- Gobet, F., & Chassy, P. (2009). Expertise and intuition: A tale of three theories. *Minds and Machines*, 19(2), 151-180.
- Hambrick, D. Z., Oswald F. L., Altmann, E. M., Meinz E. J., Gobet, F. & Campitelli, G. (2014). Deliberate practice: Is that all it takes to become an expert? *Intelligence*, 45, 34-45.
- Kaufman, S. B., & Duckworth, A. L. (2017). World-class expertise: A developmental model. *Wiley Interdisciplinary Reviews: Cognitive Science*, 8(1-2), e1365.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 363-391.
- Klein, G. (2008). Naturalistic decision making. *Human factors*, 50(3), 456-460.
- Klein, G. (2015). A naturalistic decision making perspective on studying intuitive decision making. *Journal of Applied Research in Memory and Cognition*, 4(3), 164-168.
- Krieger, L., & Bykofsky, S. (2006). *The rules of poker: Essentials for every game*. New York: Lyle Stuart.
- Laakasuo, M., Palomäki, J., & Salmela, M. (2014). Experienced poker players are emotionally stable. *Cyberpsychology, Behavior, and Social Networking*, 17(10), 668-671.
- Laakasuo, M., Palomäki, J., & Lappi, O. (2015). Pokerin psykologiaa. *Psykologia*, 50(3), 180-193.
- Laakasuo, M., Palomäki, J., & Salmela, M. (2015). Emotional and social factors influence poker decision making accuracy. *Journal of Gambling Studies*, 31(3), 933-947.
- Laakasuo, M., Palomäki, J., & Salmela, M. (2016). Poker Players With Experience and Skill Are Not ‘‘Ill’’: Exposing a Discrepancy in Measures of Problem Gambling. *Journal of Gambling and Commercial Gaming Research*, 1, 1-17.
- Leonard, C. A., & Williams, R. J. (2015). Characteristics of Good Poker Players. *Journal of Gambling Issues*, 31, 45-68.
- Levitt, S. D., & Miles, T. J. (2014). The role of skill versus luck in poker evidence from the world series of poker. *Journal of Sports Economics*, 15(1), 31-44.
- Linnet, J., Frøslev, M., Ramsgaard, S., Gebauer, L., Mouridsen, K., & Wohlert, V. (2012). Impaired probability estimation and decision-making in pathological gambling poker players. *Journal of Gambling Studies*, 28(1), 113-122.
- Linnet, J., Gebauer, L., Shaffer, H., Mouridsen, K., & Møller, A. (2010). Experienced poker players differ from inexperienced poker

- players in estimation bias and decision bias. *Journal of Gambling Issues*, 86-100.
- Macnamara, B. N., Hambrick, D. Z., & Oswald, F. L. (2014). Deliberate practice and performance in music, games, sports, education, and professions a meta-analysis. *Psychological Science*, 25(8), 1608-1618.
- Macnamara, B. N., Moreau, D., & Hambrick, D. Z. (2016). The relationship between deliberate practice and performance in sports: A meta-analysis. *Perspectives on Psychological Science*, 11(3), 333-350.
- Malmuth, M. (2012). Fundamentals of poker -- General guidelines. *TwoPlusTwo Magazine*, 8(6). Available at <http://www.twoplustwo.com/poker-fundamentals/general-guidelines/poker-etiquette.php>
- McCormack, A., & Griffiths, M. D. (2012). What differentiates professional poker players from recreational poker players? A qualitative interview study. *International Journal of Mental Health and Addiction*, 10(2), 243-257.
- Meinz, E. J., Hambrick, D. Z., Hawkins, C. B., Gillings, A. K., Meyer, B. E., & Schneider, J. L. (2012). Roles of domain knowledge and working memory capacity in components of skill in Texas Hold'Em poker. *Journal of Applied Research in Memory and Cognition*, 1(1), 34-40.
- Meyer, G., von Meduna, M., Brosowski, T., & Hayer, T. (2013). Is poker a game of skill or chance? A quasi-experimental study. *Journal of Gambling Studies*, 29(3), 535-550.
- Moreau, A., Delieuvin, J., Chabrol, H., & Chauchard, E. (2017). Online Poker Tilt Scale (OPTS): creation and validation of a tilt assessment in a French population. *International Gambling Studies*, 17(2), 205-218.
- Moreau, A., Chabrol, H., & Chauchard, E. (2016). Psychopathology of online poker players: Review of literature. *Journal of Behavioral Addictions*, 5(2), 155-168.
- O'Leary, K., & Carroll, C. (2013). The online poker sub-culture: Dialogues, interactions and networks. *Journal of Gambling Studies*, 29(4), 613-630.
- Palomäki, J., Laakasuo, M., & Salmela, M. (2013a). "Don't Worry, It's Just Poker!": Experience, Self-Rumination and Self-Reflection as Determinants of Decision-Making in On-Line Poker. *Journal of gambling studies*, 29(3), 491-505.
- Palomäki, J., Laakasuo, M., & Salmela, M. (2013b). "This is just so unfair!": A qualitative analysis of loss-induced emotions and tilting in on-line poker. *International Gambling Studies*, 13(2), 255-270.
- Palomäki, J., Laakasuo, M., & Salmela, M. (2014). Losing more by losing it: Poker experience, sensitivity to losses and tilting severity. *Journal of Gambling Studies*, 30(1), 187-200.
- Palomäki, J., Yan, J., Modic, D., & Laakasuo, M. (2016). "To bluff like a man or fold like a girl?": Gender biased deceptive behavior in online poker. *PloS one*, 11(7), e0157838.
- Parke, A., & Griffiths, M. D. (2011). Poker gambling virtual communities: The use of computer-mediated communication to develop cognitive poker gambling skills. *International Journal of Cyber Behavior, Psychology and Learning*, 1(2), 31-45.
- Rubin, J., & Watson, I. (2011). Computer poker: A review. *Artificial Intelligence*, 175(5), 958-987.
- Rhodes, M. (2010). Randy "nanonoko" Lew and his rise to multi-table poker fame. Resource document. OnlinePoker.net. <http://www.onlinepoker.net/poker-news/poker-pros-news/randy-nanonoko-lew-rise-multitable-poker-fame/7327>. Accessed October 15, 2019.
- Salen, K., & Zimmerman, E. (2004). *Rules of play : game design fundamentals* (Vol. 1st). London: MIT.
- Schlicht, E. J., Shimojo, S., Camerer, C. F., Battaglia, P., & Nakayama, K. (2010). Human wagering behavior depends on opponents' faces. *PloS one*, 5(7), e11663.
- Siler, K. (2010). Social and psychological challenges of poker. *Journal of Gambling Studies*, 26(3), 401-420.
- Sklansky, D., & Malmuth, M. (1999). *Hold'em poker for advanced players, 21st century edition*. Henderson, NV: Two Plus Two Publishing.

- Smith, G., Levere, M., & Kurtzman, R. (2009). Poker player behavior after big wins and big losses. *Management Science*, 55(9), 1547-1555.
- Talberg, N. (2019). Learning Poker in Different Communities of Practice: A Qualitative Analysis of Poker Players' Learning Processes and the Norms in Different Learning Communities. *Journal of Gambling Issues*, 42, 8-41.
- Taylor, I., & Hilger, M. (2007). *The poker mindset: Essential attitudes for poker success*. Dimat Enterprises Incorporated.
- Tendler, J. (2011). *The mental game of poker: Proven strategies for improving tilt control, confidence, motivation, coping with variance, and more*. Self-published, theScore esports (2019, March 24). What is Tilt? The Gentle Art of Not Losing Your Sh*t in Esports [Video file]. Retrieved from <https://www.youtube.com/watch?v=OvaPVB FVUWQ>.
- Thompson, J. J., McColeman, C. M., Stepanova, E. R., & Blair, M. R. (2017). Using video game telemetry data to research motor chunking, action latencies, and complex cognitive-motor skill learning. *Topics in Cognitive Science*, 9(2), 467-484.
- van Loon, R. J. P., van den Assem, M. J., & van Dolder, D. (2015). Beyond chance? The persistence of performance in online poker. *PloS one*, 10(3), e0115479.
- von Neumann, J., & Morgenstern, O. (1944/2007). *Theory of games and economic behavior*. Princeton: Princeton University Press.
- Wright, R. (2001) *Nonzero: The logic of human destiny*. New York: Vintage.
- Wolkomir M. (2012). "You fold like a little girl!" (Hetero) Gender framing and competitive strategies of men and women in no limit Texas hold em poker games. *Qualitative Sociology*, 35 (4), 407-426.

Received: 16 October 2019

Revision received: 04 March 2020

Accepted: 10 March 2020



Appendix

Task Analysis of a Poker Decision

This analysis of the most popular game variant of poker, namely No Limit Texas Hold’Em (NLHE), 1) outlines the game’s rules and strategic fundamentals, and 2) illustrates the cognitive complexity of the game.

The goal in NLHE, like any other poker variant, is to have the best combination of cards in comparison to the other players on the table. In NLHE the sequence in one round of play goes as follows. First, during so-called “pre-flop,” two cards are dealt to each player (“hole cards”); these cards can be seen only by the respective player. Next, five commonly shared cards (“community cards”) are sequentially dealt to the middle of the table for everybody to see. The first three community cards are known as “the flop” and dealt simultaneously. The following two community cards are known as the “turn” and “river,” and dealt separately. These community cards in combination with the hole cards determine the winner of the round. To be more precise, the players can use one or both (or none) of their hole cards in combination with the community cards to form the best five-card combination, following a hierarchical ordering of values of all five-card hands.

There are, at most, four rounds of betting: (1) during pre-flop, (2) after the flop is dealt, (3) after the turn is dealt, and (4) after the river is dealt. The betting has the following structure. Players can either “check” (not bet while not giving up), “bet” (invest money in the pot), “call” (match an opponent’s bet), “raise” (go beyond an opponent’s bet), or “fold” (give up and exit the round). Once a full round has been played this is considered as having played one “hand.”

To illustrate the logic of the game let us examine the situation outlined in Figure A.1. Your (“YOU” in Figure A.1) current best five-card hand is called *king high*, which is considered a very weak hand. However, one more community card (the *river*) will be dealt if neither player still in contention gives up (*folds*). Hence, there is another chance for you to improve your hand. One way to improve is when the river card is either a “4” or a “9”; then you would have a *straight*, which is the best possible hand given the

current community cards. Deciding whether to stay in the game, or not, is influenced by the likelihood of this event occurring. In this case, the odds of your hand improving into a straight on the river is about 17% (at maximum 8 cards from a total of 42).

Let us assume that you are highly skilled and decide to pass the turn to the Opponent (*check*). The Opponent bets \$100 into the pot of \$135 (making the pot \$235 in total). Now you know that to be guaranteed to win (disregarding ties for simplicity) your hand needs to improve into a straight on the river. Matching the opponent’s bet (*calling*) of \$100 when the size of the pot is \$235 increases the pot to \$335. This corresponds to *immediate odds* of 2.35 to 1 (or $100/335 = 29.8\%$), which means that calling would be profitable if it were the winning decision 29.8% of the time. In simpler terms, one would need to be in a similar situation at least 2.35 times for the same decision to have a positive outcome.

Since your hand will improve into a straight on the river only about 17% of the time, it follows you should *not* call based on the *immediate odds* alone ($17\% < 29.8\%$). However, you also know that *if* you call *and* improve your hand into a straight, you might win additional money by making the pot larger—either by betting yourself or by “inducing” a bet from the Opponent by checking. Thus, you should consider also your so-called *implied odds*; that is, what calling now might imply later in terms of profit.

Whether the implied odds justify calling depends on the Opponent’s strategy and the hand the Opponent is holding. For instance, if the Opponent is unskilled it might be rational to take the chance of playing despite the poor immediate odds, because unskilled players are more likely to make mistakes and “pay off” bets on the river when they should not. In other words, even if a certain card combination would be statistically unlikely to win, in certain situations it might still make sense to play them.

In poker, players have only probabilistic information on how to act and need to rely on previous experience and reasoning skills to make their next decisions. This process involves estimating all of the possible card combination the Opponent is expected to have (“hand range”),

given the community cards and the Opponent's betting actions previously (and body language in live poker; or chat comments in online poker, and so on).

In the situation outlined in Figure A.1, you can estimate how "strong" your own five-card hand is against the "average strength" of the Opponent's hand range. This estimation is sometimes done quickly and implicitly, because time pressure alone often prevents explicit detailed calculations – skilled players sometimes play on multiple tables online, some on as many as 24 at a time (e.g., Rhodes, 2010).

The analysis above is an oversimplification, and merely illustrates the complexities in poker decision-making. You as a player in the game should also consider *bluffing* on the river, if your hand does not improve. Also, you could decide to *bet* initially, or *raise* the Opponent's initial bet after checking. These would entail new probabilistic dependencies, which we have omitted. While this task analysis is hypothetical, it is an empirical question how explicitly analytical (or intuitive) players' cognitive processes are in similar situations. Determining opponents' hand ranges in various situations is discussed across poker communities (O'Leary & Carroll, 2013).



Figure A.1. An online NLHE table (adapted from Palomäki et al., 2016). A: Opponents 1, 3 and 4 have folded (given up) and are no longer contesting the pot. B: The total amount currently in the pot, which represents all the money that has been previously waged during the current hand. C: The amount of money the player has remaining in their stack, which represent the total amount they will be able to wage during any particular hand. D: The "hole cards" of the Player, not visible to the opponents. E: The "hole cards" of the remaining opponent. F-H: The "community cards" shared by the player and the opponent. F: The flop (three first "community" cards). G: The turn (the fourth community card). H: The river (the fifth and last community card to be dealt, at this point unknown).