

Skewness Preferences: Evidence from Online Poker

Markus Dertwinkel-Kalt, Johannes Kasinger, Dmitrij Schneider

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Skewness Preferences: Evidence from Online Poker

Abstract

We test for skewness preferences in a large set of observational panel data on online poker games ($n=4,450,585$). Each observation refers to a choice between a safe option and a binary risk of winning or losing the game. Our setting offers a real-world choice situation with substantial incentives where probability distributions are simple, transparent, and known to the decision-makers. Individuals reveal a strong and robust preference for skewness, which is inconsistent with expected utility theory. The effect of skewness is most pronounced among experienced and unsuccessful players but remains significant in all subsamples that we investigate, in contrast to the effect of variance.

JEL-Codes: D010, D810, G400.

Keywords: risk preferences, choice under risk, skewness, gambling.

Markus Dertwinkel-Kalt
University of Münster / Germany
markus.dertwinkel-kalt@wiwi.uni-muenster.de

Johannes Kasinger
Tilburg School of Economics and Management
Tilburg University / The Netherlands
j.kasinger@tilburguniversity.edu

Dmitrij Schneider
Düsseldorf Institute for Competition
Economics (DICE), Düsseldorf University
Düsseldorf / Germany
schneider@dice.hhu.de

February 2024

We thank Björn Bartling, Sabine Bernard, Thomas Dohmen, Markus Eyting, Katrin Gödker, Max R. P. Grossmann, Andreas Hackethal, Matthias Heinz, Simon Heß, Michael Kosfeld, Mats Köster, Jan Krahnert, Leon Landes, Christine Laudenbach, Maximilian Schmetzke, Christian Seel, Matthias Sutter, and Martijn van den Assem for constructive comments and suggestions. We further thank seminar participants at the Goethe University Frankfurt, the Heinrich Heine University Düsseldorf, and the Leibniz Institute for Financial Research SAFE as well as several conference audiences for valuable feedback. We gratefully acknowledge financial support by the German Research Foundation (DFG project No. 462020252).

1 Introduction

Over the last 50 years, scholars in finance and behavioral economics have argued that whether people take up or avoid risk depends on its skewness, that is, on its standardized third moment. *Skewness preferences*—a preference for positive and an aversion toward negative skewness—are a central prediction of most behavioral models of choice under risk, such as prospect theory (Kahneman and Tversky, 1979), salience theory (Bordalo *et al.*, 2012), or models building on the anticipation of regret or disappointment (Loomes and Sugden, 1982, 1987a,b; Gul, 1991; Inman *et al.*, 1997). In contrast, expected utility theory cannot explain why people seek positively skewed risks but avoid negatively skewed risks. Skewness preferences can explain a wide range of seemingly disparate puzzles in choice under risk: the favorite-longshot bias, whereby positively skewed long shots are overbet and negatively skewed return distributions of favorites are underbet (e.g., Snowberg and Wolfers, 2010); the simultaneous demand for lottery-like gambles and high-premium insurances (Kahneman and Tversky, 1979; Sydnor, 2010; Garrett and Sobel, 1999); the Allais paradox (Allais, 1953); the growth puzzle (Fama and French, 1992; Bordalo *et al.*, 2013); and many instances of portfolio underdiversification (Mitton and Vorkink, 2007).

Despite the vast literature on skewness preferences (for a survey see Trautmann and van de Kuilen, 2018), direct tests of the effect of skewness on real-world choices are scarce: the support for skewness preferences either comes from highly stylized laboratory experiments (e.g., Ebert and Wiesen, 2011; Ebert, 2015; Dertwinkel-Kalt and Köster, 2020), or—mainly in finance and labor economics—from real-world decision situations where the underlying probability distribution is very complex, estimated based on past data, and basically unknown to the decision maker.¹

This paper directly tests for skewness preferences in a large set of observational panel data on real-world choices. Our setting combines a number of advantages. Contrary to laboratory experiments, our study covers a wide range of incentives, with an expected value of the average lottery of \$62.39, and builds on a high number of observations and less artificial choice situations. In contrast to the empirical finance and labor economics literature, we do not have to approximate the underlying probability distribution and do not need to impose strong assumptions that individuals correctly estimate and comprehend such a distribution, as probabilities are transparently displayed to decision-makers before they make their decisions. The risks involved are binary and uniquely determined by the first three moments of their probability distribution: expected value, variance, and skewness, which allows for a clean identification of skewness preferences. In addition, individuals in our data set can generally be expected to understand risks and probabilities well, reducing the confounding effects of misunderstandings of probabilities and imperfect information.

In detail, we study risk *preferences* in online poker, making use of a novelty that one of the leading online poker platforms, *Pokerstars*, introduced in August 2019: the so-called *all-in cashout*. The all-in cashout provides insurance against a player's risk in a showdown situation.

¹Among others, the tests of skewness preferences in finance (Boyer *et al.*, 2010; Bali *et al.*, 2011; Conrad *et al.*, 2013; Lin and Liu, 2018; Jondeau *et al.*, 2019) and labor economics (Hartog and Vijverberg, 2007; Berkhout *et al.*, 2010) fall into this category: here, the probability distributions can only be approximated through expected moments that are estimated based on past data.

In a showdown situation, the outcome of the poker hand is solely determined by the cards drawn from the remaining deck of cards. The two possible outcomes for each player are: i) losing and receiving a payout of zero or ii) winning the entire *pot*, that is, the accumulated bets by players throughout a hand. The all-in cashout now gives each player in a showdown situation the additional option to choose a safe payout equal to the expected payout of the underlying lottery minus a profit margin for *Pokerstars* of 1%. Before making her decision whether to take the insurance or not, all relevant information is disclosed to each player. Each observation in our data set refers to a player's choice between this safe insurance option and the respective binary lottery.² Our data set includes 4,450,585 of such individual choices in showdown situations with two opponents for Omaha Poker cash games, all collected between January 1, 2020, and June 30, 2021.³

This specific setting allows us to identify the effect of skewness on risk-taking. We show that skewness has a sizeable effect on risk-taking, while the effect of variance is comparably negligible. The insurance option is selected in 20.0% of cases when the risk is left-skewed, but only in 14.2% of cases when the risk is right-skewed. Put differently, it is around 40% more likely that the insurance option is selected when players face a left-skewed instead of a right-skewed risk. As winning the pot is a complementary event for the two opponents, for each observation involving a right-skewed lottery with the winning probability $\pi < 0.5$, there is exactly one observation in our data set that involves a left-skewed lottery with the winning probability $1 - \pi > 0.5$. Both of these lotteries have identical variance but inverse skewness, which allows us to circumvent the limitations of other studies (e.g., Golec and Tamarkin, 1998) where variance and skewness always change simultaneously.

In our regression analyses, we follow Mitton and Vorkink (2007) and assume a utility function that is linear in the different risk moments. In our basic specification, we regress an insurance choice dummy that equals one if player i chooses the insurance option in showdown j and zero otherwise, on the first three moments of the underlying lottery. We further include game fixed effects to control for heterogeneity in insurance choices across different games and month fixed effects to account for month-specific heterogeneity as potentially driven by seasonal effects or COVID-19 countermeasures. Increasing skewness by one standard deviation, keeping the variance and expected values of the lotteries constant, decreases the likelihood that the insurance option is chosen by about 2.3 percentage points, which is equivalent to a decrease of around 13.5% compared to the average share of positive insurance choices (i.e., the mean dependent variable). The estimated effect of variance is negligible in most specifications. A one standard deviation increase in variance is associated with a decrease in insurance choice of around 0.0 to 0.3 percentage points, depending on the specification. In the basic linear probability model, the estimated effect of variance is statistically insignificant.

These results remain robust to different empirical specifications (such as Probit and Logit models), to controlling for player-specific characteristics (such as experience and average profit

²Strictly speaking, the lottery is not always binary as there are situations where a split pot can occur; we discuss this limitation of our study in Section 5.1.

³Showdowns with more than two players typically yield more complex probability distributions, depending on the timing of bets during a hand and on players' available budgets. Thus, for clarity, we restrict our analysis to two-person showdowns.

per hand) and hand-specific variables (such as the amount of money the player started the hand with, the weekday, or the *stake*, that is, the size of mandatory bets), to excluding outliers, and to using the coefficient of variation (the inverse "Sharpe" ratio of the lottery) instead of expected value and variance. The panel structure of our data further allows us to include individual fixed effects to control for time-invariant heterogeneity across individuals. Including individual fixed effects does not change the impact of skewness, neither in magnitude nor significance, and it also does not substantially affect the estimated coefficients of the first two moments.

We also explore whether our effects vary across different subgroups. To determine whether our findings hold in contexts with large monetary stakes, we restricted our analysis to the subsample that involves pots of more than \$100. In this subsample, the average lottery has an expected value of \$252.93. The proportion of insurance choices (11.7%) is notably lower than in the full sample. Despite this, the skewness effect's magnitude is only slightly smaller, which suggests that the effect of skewness on risk-taking behavior is similar across both low- and high-stake decisions.

Moreover, we split our sample at the median for various player- and hand-specific characteristics. We find evidence of skewness preferences in all subsamples. The evidence is strongest for relatively experienced players, as measured by the number of showdowns the player participated in or by the total number of hands played (including hands without showdowns). For this group of players, the estimated effect of skewness is about one and a half times the effect in the full sample. For this group, the effects of the other two moments are either smaller in absolute terms or even change signs compared to the full sample. Additionally, we analyze the impact of skill by considering separate analyses for successful (those with positive aggregated net profits over our observation period) and unsuccessful players (those with aggregated net losses). While a player's skill cannot affect the outcome of the studied lotteries, it can affect the long-run profits of players and is likely related to differences in risk attitudes and players' motivation to play poker. Indeed, while skewness affects both successful and unsuccessful players, its effect is stronger for the latter group.

Finally, we disentangle skewness effects from loss aversion. Players could perceive their contribution to the pot as a reference point. According to loss aversion, they exhibit a preference to at least recoup this reference amount, thus achieving a gain and avoiding a loss. In our setting, players facing left-skewed risks usually recoup their contribution with the insurance option and thus realize a gain. Conversely, players who experience a loss from choosing the insurance option may prefer to take the risk as this offers the chance of a gain. Players facing right-skewed risks usually fall into this category. There are exceptions, however, because of "dead" money from players not reaching the showdown. These instances of players choosing between a right-skewed risk and an insurance option that offers a gain compared to their pot contribution allow us to disentangle loss aversion and skewness preferences. We find evidence for both skewness preferences and loss aversion. When the insurance option represents a gain instead of a loss, players are significantly more likely to choose the insurance (18.2% vs 13.3%). When controlling for loss aversion in our regression analyses, the magnitude of the skewness coefficient is roughly half but remains highly significant. We thus find evidence for the coexistence of skewness preferences and loss aversion.

Among others, this paper contributes to the literature on decision-making in poker games (for other related literature, see Section 2). Besides some studies that aim to quantify the extent of skill and luck in poker games (e.g., Fiedler and Rock, 2009; Potter van Loon *et al.*, 2015; Duersch *et al.*, 2020), researchers have mainly used poker data to study reference-dependent risk attitudes. Smith *et al.* (2009) and Eil and Lien (2014) find that poker players play less cautiously, longer, and more aggressively after losing a big pot or if a player is losing within a poker session. These findings are in line with the break-even hypothesis predicted by loss aversion (Kahneman and Tversky, 1979; Thaler and Johnson, 1990). One challenge when analyzing poker play is that it crucially depends on the players' expectations about the opponents' hands and playing style (both of which are unobservable). By focusing on those situations where all uncertainty is resolved and players cannot actively affect the outcome of the game anymore, we can, however, overcome these problems. So unlike in previous studies (e.g., Smith *et al.*, 2009; Eil and Lien, 2014), missing information, wrong beliefs or a misunderstanding of the risks involved should not play a role in our setup and should not confound our insights on the drivers of risk-taking.

We proceed as follows. In Section 2 we define skewness preferences and discuss the related literature. In Section 3 we describe our data and setting before we present our results in Section 4. Section 5 discusses the limitations of our framework and provides corresponding robustness checks. Section 6 concludes.

2 Theoretical Background and Related Literature

Our observations involve binary decisions between a safe "insurance" option and a binary lottery L . As we will see in Lemma 1, such a binary lottery is uniquely defined by its expected value $\mathbb{E}[L]$, its variance $Var[L]$ and its skewness $S[L]$, which is defined by the third standardized central moment

$$S[L] := \mathbb{E} \left[\left(\frac{L - \mathbb{E}[L]}{\sqrt{Var[L]}} \right)^3 \right]. \quad (1)$$

We can then define the following notions.

Definition 1. Lottery L is called right-skewed (or, equivalently, positively skewed) if $S(L) > 0$, left-skewed (or, equivalently, negatively skewed) if $S(L) < 0$, and symmetric otherwise.

Other definitions of positive skewness, such as via "long and lean" tails of the risk's probability distribution, exist and are, in general, not equivalent. For binary risks $L = (x_1, \pi; x_2, 1 - \pi)$, where outcome x_1 is realized with probability π and x_2 with probability $1 - \pi$, Ebert (2015), however, shows that all conventional notions of skewness are equivalent and the skewness of a binary risk is well-defined. Moreover, binary lotteries can be uniquely identified by their first three moments as shown by Ebert (2015) and generalized by Dertwinkel-Kalt *et al.* (2023):

Lemma 1. For constants $E \in \mathbb{R}$, $V \in \mathbb{R}_+$ and $S \in \mathbb{R}$, there exists exactly one binary lottery $L = (x_1, \pi; x_2, 1 - \pi)$ with $x_2 > x_1$ such that $\mathbb{E}(L) = E$, $Var(L) = V$ and $S(L) = S$. Its parameters are

given by

$$x_1 = E - \sqrt{\frac{V(1-\pi)}{\pi}}, \quad x_2 = E + \sqrt{\frac{V\pi}{1-\pi}}, \quad \text{and } \pi = \frac{1}{2} + \frac{S}{2\sqrt{4+S^2}}. \quad (2)$$

We denote the binary lottery with expected value E , variance V , and skewness S as $L(E, V, S)$.

As a result of Lemma 1, it is possible to vary skewness for binary risks while fixing the first two moments. We can now define skewness preferences as follows:

Definition 2. *An agent reveals a preference for skewness if the following holds: for any $E \in \mathbb{R}$ and any $V \in \mathbb{R}_+$, there exists a unique threshold value $\hat{S} \geq 0$ so that she strictly prefers the binary lottery $L(E, V, S)$ over the safe option that pays E if and only if $S > \hat{S}$.*

Expected Utility Theory. In the following, we formally show that Definition 2 violates expected utility theory (henceforth: EUT) and is not coherent with *any* EUT utility function. Let $u(\cdot)$ be a utility function with normalization $u(0) = 0$, and consider $L = (0, 1 - \pi; x, \pi)$. According to Definition 2, there exists a threshold skewness value, or equivalently, a threshold probability level $\bar{\pi} \in (\frac{1}{2}, 1)$, so that for all $x \geq 0$ and all $\pi \in (0, 1)$ we have

$$u(\pi x) \geq \pi u(x) \text{ if } \pi > \bar{\pi} \quad (3)$$

$$u(\pi x) < \pi u(x) \text{ if } \pi \leq \bar{\pi}. \quad (4)$$

Now we define $x' := 2x$ and assume $\pi \in (\frac{1}{2}, \frac{\bar{\pi}}{2})$. Then, (3) and (4) yield $2\pi u(x) \leq u(2\pi x) = u(\pi x') < \pi u(2x)$, and thus $2u(x) < u(2x)$. But we also have $u(x) = u(\frac{1}{2} \cdot 2 \cdot x) \geq \frac{1}{2}u(2x)$, which gives $2u(x) \geq u(2x)$. Taken together, these two derivations give a contradiction. Intuitively, Definition 2 stipulates that the agent is (i) risk-averse over all risks that are not sufficiently skewed, which necessitates a weakly concave utility function, but (ii) risk-seeking over all sufficiently skewed risks, which necessitates a strictly convex utility function. In sum, (i) and (ii) give a contradiction.

Notably, in our empirical setting the agent cannot obtain a gamble's expected value E , but only slightly less, namely $0.99E$. This, however, does not alter any of the conclusions we have drawn on the validity of EUT, as demonstrated in the following. Suppose an agent strictly prefers a risky option $L = (0, 1 - \pi; x, \pi)$ over 99% of its expected value, πx , if and only if the risky option's skewness is sufficiently large. Then we obtain the conditions that are analogous to (3) and (4), namely,

$$u(0.99\pi x) \geq \pi u(x) \text{ if } \pi > \bar{\pi}$$

$$u(0.99\pi x) < \pi u(x) \text{ if } \pi \leq \bar{\pi}.$$

The very same contradiction as above can be constructed from these conditions. Intuitively, in order to weakly prefer $0.99E$ over a risky option with expected value E and non-positive skewness, the utility function must be sufficiently concave, but in order to prefer the risky option

over $0.99E$ when skewness is sufficiently positive, the utility function must not be that concave—a contradiction.

Contrary to our definition, some papers in the literature, such as Ebert and Karehnke (2020), define skewness preferences as preferring lottery $L(E, V, S)$ over $L(E, V, -S)$. Skewness preferences in this alternative sense can be reconciled with EUT by making the additional assumption of prudence (whereby the third derivative of the utility function is strictly positive). We, however, view this alternative definition as less relevant for practice. In fact, we are unaware of any real-world choice between $L(E, V, S)$ and $L(E, V, -S)$, where an agent chooses between two risks that only differ in the sign of skewness. All of the motivating examples in the first paragraph of our introduction (e.g., gambling vs. insurance choice, and most of the applications in financial or labor economics), as well as the choice situations that we investigate in this paper, better fit our definition where individuals choose between a *risky* and a *safe* option. When an individual decides whether to gamble or not, she chooses between some risk (gambling) and the safe option of not gambling. When an individual decides whether to buy insurance, she chooses between the risky option (of not taking up insurance) and the safe option of buying insurance. Explaining these choices between a safe and a risky option is also the core contribution of prospect theory, as highlighted in the last sentence of the abstract of the seminal paper by Kahneman and Tversky (1979).

Taken together, skewness preferences as defined in Definition 2 are not coherent with EUT. In order to explain why agents dislike symmetric risks, EUT needs to assume that the utility function is strictly concave. Under this assumption, however, EUT predicts, for any skewness level, that the safe payout of E should be strictly preferred over any lottery that pays E in expectation. When EUT wants to explain why an agent takes up a positively skewed risk, it has to assume a convex utility function ($u'' > 0$), but this then goes along with the implausible prediction (that also violates Definition 2) that *every* binary lottery is preferred over the safe option that pays its expected value. So, EUT cannot explain why people's preference to take up a risk depends on the risk's skewness. The same holds for standard portfolio theory (Markowitz, 1952), whereby people's utility from some risk is linearly increasing in its expected value and linearly decreasing in its variance. Thus, EUT and standard portfolio theory would predict that whether some risk is taken up mainly depends on variance, but not on skewness.

Behavioral Economics. Skewness preferences as defined in Definition 2 are predicted, however, by most behavioral models of choice under risk such as cumulative prospect theory (Kahneman and Tversky, 1979), salience theory (Bordalo *et al.*, 2012), regret theory (e.g., Bell, 1982; Loomes and Sugden, 1982), and disappointment aversion (Gul, 1991), as shown, for instance, in Barberis (2012), Dertwinkel-Kalt and Köster (2020), and Ebert and Karehnke (2020). Also seminal models proposed in the behavioral finance literature (Mitton and Vorkink, 2007) allow for skewness preferences by, for instance, augmenting standard portfolio theory by an additional term that allows not only for a linear effect of expected value and variance, but also of skewness on utility.

Skewness preferences allow us to understand why revealed attitudes toward risks vary across contexts. On the one hand, people like to gamble (e.g., Golec and Tamarkin, 1998,

Garrett and Sobel, 1999) but they also overpay for insurance with low deductibles (e.g., Sydnor, 2010; Barseghyan *et al.*, 2013). In financial markets, investors seek positively skewed return distributions (Chunhachinda *et al.*, 1997; Prakash *et al.*, 2003; Mitton and Vorkink, 2007; Boyer *et al.*, 2010; Bali *et al.*, 2011; Conrad *et al.*, 2013). Skewness preferences also matter in labor economics and, in particular, career choices (Hartog and Vijverberg, 2007; Berkhout *et al.*, 2010; Grove *et al.*, 2021) as workers accept a lower expected wage if the distribution of wages in a cluster (i.e., education-occupation combination) is right-skewed. Similarly, a preference for skewness can explain the substantial entrepreneurial investments in private equity with bad risk-return tradeoffs (Moskowitz and Vissing-Jørgensen, 2002).

Implications for our empirical approach. The characteristics of binary lotteries outlined above make it appealing to study skewness effects at the hand of binary lotteries. To identify skewness preferences, it is optimal to let agents repeatedly choose between a safe option and a binary lottery, where only the lottery’s skewness differs between choices. In such a setting, skewness preferences predict a negative relation between insurance choice and the lottery’s skewness. While this decision situation is hardly implementable in the field (for a laboratory experiment that precisely implements this see Experiment 1 in Dertwinkel-Kalt and Köster, 2020), our setup approximates these experiments as closely as possible (for a discussion of differences to the ideal setup see the discussion of limitations in Section 5.1).

In our setup, poker players face the choice between a lottery and the safe option that pays 99% of the lottery’s expected value. Thus, the insurance is selected if and only if for the player’s utility function $U(\cdot)$ we have $0.99 U(\mathbb{E}(L)) > U(L)$. We follow Mitton and Vorkink’s (2007) reduced-form approach by assuming "Lotto investors" that have identical preferences as traditional investors over mean and variance (see Markowitz, 1952), but also a preference for skewness. The utility such investors derive from some lottery L is then given by

$$U(L) = \mathbb{E}(L) + \beta_V \text{Var}(L) + \beta_S S(L).^4$$

We will not directly estimate the effect of the lottery’s moments on utility, but on the likelihood that the safe insurance option is preferred over the lottery. Given this dependent variable, a positive (negative) coefficient β_V indicates variance-averse (variance-seeking) agents, and a positive (negative) coefficient β_S indicates a preference for negative (positive) skewness. This approach, therefore, allows for both positive and negative effects of variance and skewness on insurance choice.

3 Background and Data

In this section, we first provide background information about the underlying poker game and the new insurance option that our study builds on (Section 3.1). In Section 3.2, we give an overview of our data set.

⁴Unlike Mitton and Vorkink (2007) we adopt the usual narrow-framing assumption that is adopted throughout experimental economics: namely, that subjects do not integrate their earning from the respective game into their overall wealth, but evaluate it in isolation.

3.1 Background

On Omaha Poker cash games. We analyze hands from Omaha Poker cash games. In a cash game, all players start the hand with an amount of real money, the *stack*, which will be used for betting throughout the respective hand.⁵ Money cannot be added or withdrawn during a hand. However, players can leave the game after a hand is concluded or add chips up to a maximum amount depending on the blinds (i.e., the mandatory bets posted before every hand) of the respective game. Accordingly, we define the stake of a game by the size of the blinds. In a poker cash game there are usually two blinds, the *big blind* and the *small blind*, which is half the size of the big blind. In the remainder of the paper, the stake always refers to the big blind.

In Omaha Poker cash games, each player is dealt four private cards (*hole cards*) that are only visible to the respective player. In addition, there are up to five *community cards* that are public information and are dealt throughout three stages: i) *Flop*: first three community cards; ii) *Turn*: fourth community card; iii) *River*: fifth community card. Each stage is preceded and/or followed by a betting round. The money that players post throughout these betting rounds is collected in the *pot*. Furthermore, there is a fee collected by Pokerstars, called the *rake* that is deducted from the pot. The rake is calculated as a percentage of the pot, ranging from 3.5% to 5% depending on the stake, and capped at a certain amount. Accordingly, the winning player is awarded the *net pot*, that is, the pot minus the rake.

If the betting causes all but one player to lay down their hole cards (i.e., they *fold*), the remaining active player wins the net pot without showing any private cards. Otherwise, the net pot is awarded to the active player with the best five-card poker hand after the last community card is dealt. This best five-card poker hand consists of two of the player's hole cards and three community cards (see also the official ranking of poker hands in Appendix A.1).

The hole cards are revealed when either the betting round after the River is finished or when there are $N > 2$ active players, of which at least $N - 1$ players are *all-in*, that is, they put their entire stack in the pot. The latter scenario is called a *showdown*. In a showdown, the players face a binary lottery L , whose outcome depends solely on the cards that will be drawn from the remaining deck of cards: $L = (\text{pot} - \text{rake}, \pi; 0, 1 - \pi)$. In such a situation, the probability that one player wins the net pot, π , can be calculated conditional on the revealed individual hole cards and the community cards that have been dealt until the showdown. Under certain circumstances, split pot situations may occur. In such a scenario, the lottery is not binary. We abstract from these scenarios here and discuss the issue in Section 5.1 in more detail.

Figure 1 shows an example of a showdown after the Turn, that is, with one card to come. Player 1 is all-in, and no more betting is possible. The best five-card hand of Player 1 is a *High Card Queen*, which loses against the *Two Pairs* (queens and tens) of Player 6. Player 1 can only win the hand if a *Heart*-card is drawn from the remaining cards, which would give her a winning *Flush* (five cards of the same suit). In total, there are 13 Heart-cards in the 52-card deck. Five

⁵Beside the fact that the insurance option is only available for cash games, cash games are also more suited to our question and "easier to analyze than tournament games, since in a cash game, a player who is risk neutral over money should also be risk neutral over chips. This is not necessarily the case in a tournament, for a number of reasons" (Eil and Lien, 2014), including varying incentives to outlast other players in different tournament phases or non-linear payout structures.

Figure 1: Exemplary screen of a showdown on Pokerstars with one card to come



Heart cards have been already revealed, implying that there are still eight Heart cards among the 40 cards that have not been revealed yet. The probability of Player 1 winning the hand is thus simply the number of remaining Heart cards (eight in our example) divided by the number of remaining cards, $52 - 12 = 40$, so that we obtain a winning probability of Player 1 of $\frac{8}{40} = 0.20$, and a probability of losing of 0.80. As shown in Figure 1, these probabilities and the exact size of the net pot are displayed in an all-in situation on the players' screens. If the showdown happens at an earlier stage, the probabilities can be calculated by dividing all possible realizations of community cards, in which a specific player holds the winning hand, by the total number of possible realizations. Again, the corresponding probabilities and the net pot are displayed to the players (see in Appendix A.2 an example where the showdown happens before any community card is revealed).

The insurance option. We make use of the new insurance option (the so-called *all-in-cashout*) introduced on August 13, 2019, on the Pokerstars website, which provides a safe alternative against the risk that the players face in a showdown. If a player chooses this insurance option, she will no longer be eligible to contest any portion of the net pot, and the offered amount will be added to her stack immediately and risk-free. If she declines, she will continue to contest the entire net pot as usual. Importantly, while other players in a game can observe the other players' insurance decisions, the outcome of a showdown is unaffected by the opponent's insurance

decision. Players declining the insurance option still need the best hand in a showdown to win the net pot, even if all their opponents have cashed out. As a result, each active player in a showdown faces a choice between a binary risk and a safe option. The guaranteed payout from choosing the insurance option is equal to the expected value of the lottery minus a fee of 1% on this expected value, that is, equal to $$(pot - rake) \times \pi \times 0.99$. The 1% fee charged by Pokerstars is equal for all players and has not changed since the insurance option's introduction. The insurance payout is rounded to full cents.

To better understand a player's decision in a showdown, turn again to Figure 1, which shows a situation in which the insurance option is offered to Player 1. The two red buttons represent the binary choice between two options: i) the risky option "Resume hand" (explained above) that pays \$2.79 with 20% and zero with 80%; and ii) the safe insurance option that pays \$0.55 with 100%. As can be seen in Figure 1, the probabilities and the exact size of the net pot are displayed on the players' screens. The displayed insurance payout already includes the rake and the 1% fee. Accordingly, Player 1's insurance payout in our example is equal to: $\$2.79 \times 0.2 \times 0.99 = \0.55 and Player 6's insurance payout will be: $\$2.79 \times 0.8 \times 0.99 = \2.21 . Thus, the players are readily presented with all information that is relevant for their decision. The players have 12 seconds to make their choice. If players do not choose one of the proposed alternatives within 12 seconds, the hand resumes with the risky option.

3.2 Data

Our data set includes 4,450,585 observations, where every observation refers to a unique decision by a single player in a two-person showdown situation as described above. This includes the decisions of 83,219 distinct players.⁶

Our data set is extracted from 35,529,631 distinct poker hands played between January 01, 2020, and June 30, 2021 on Pokerstars, the largest online poker network during our observation period (Primedope, 2023). The number of distinct poker hands exceeds the number of observations, as not every poker hand results in a two-person showdown. We obtain the raw data from a commercial poker data provider that collects and stores hand histories for every Omaha Poker cash game played on Pokerstars.⁷ Hand histories are automatically generated by the Pokerstars software and include all public information of a single poker hand. The raw data is provided in text files and we make use of the commercial poker software "PokerTracker 4" to convert the raw data into a workable data set. Smith *et al.* (2009) use an earlier version of "PokerTracker" to construct their data set.

Our data set spans a wide range of games that vary in the size of the mandatory bets and the maximum amount players can bring to the respective game. The hand histories include information on all community cards, the hole cards of the players that went to a showdown, and

⁶As we explain in Section 5.1 in more detail, we exclude all observations that involve a winning probability of 0.5, as most of these observations refer to situations with degenerated lotteries, for which no insurance option was offered.

⁷Our data provider *HH Dealer* (<https://www.hhdealer.com/>) collects the hand histories for various online poker platforms. At Pokerstars, a hand history is accessible by all Pokerstars users that opened the window of a respective game. As a quality check, we observed several sessions on the Pokerstars platform during the data collection period and checked whether the obtained histories are complete and accurate.

the net pot size. This information allows us to calculate each player’s probability of winning the pot at showdown and assign each player’s insurance payout value. Furthermore, our data allow us to infer whether a respective player has chosen the insurance option for each showdown.

We capture a player’s decision between the safe option and the binary lottery by the variable *insurance choice*, our dependent variable of interest. This variable equals one if the player chooses the safe option and zero if she chooses the binary lottery. In our data set, players choose the safe option in 17.1% of cases with a standard deviation of 0.377 (Table 1). Notably, the overall share of choices of the insurance option is rather low. A majority of players choose the risk instead of the insurance option, which might be driven by a large share of individuals that are risk-seeking in general or by the 1%-margin that has to be paid to *Pokerstars* if the insurance option is chosen.

Table 1: Summary statistics on insurance choice

Statistic	N	#(Choice=1)	#(Choice=0)	Mean	St. Dev.
Insurance Choice	4,450,585	761,585	3,689,000	0.171	0.377

Note: The table reports summary statistics on the insurance choice dummy that equals one if the safe option is chosen and zero otherwise.

The winning probabilities and the net pot size allow us to calculate the expected value, $E = \pi x$, the variance, $V = \pi(1 - \pi)x^2$, and the skewness, $S = \frac{1-2\pi}{\sqrt{\pi(1-\pi)}}$, of each binary gamble players face in a showdown situation. Table 2 presents the descriptive statistics for the first three moments of the binary lotteries in our data set. Our lotteries have an average expected value of 62.39, with a standard deviation of 326.58 and a median of 13.37 (all in US-\$). The expected values range from as little as \$0.001 to as much as \$64,242.40. Our measure of skewness takes values between -40.79 and 40.79 with a mean of zero.

Table 2: Summary statistics on different moments

Statistic	N	Mean	St. Dev.	Pctl(25)	Median	Pctl(75)
Expected Value	4,450,585	62.39	326.58	5.11	13.37	35.92
Variance	4,450,585	73,786.68	2,157,525.00	29.52	148.86	964.71
Skewness	4,450,585	0.00	2.23	-0.86	0.00	0.86

Note: The table reports summary statistics of the expected value, variance, and skewness for the lotteries in our sample, using net pot sizes measured in US-\$.

The data set also includes a pseudonymized player ID, which serves as a unique identifier for each observation (together with the distinct hand number). In addition, our data records each player’s stack at the beginning of the hand and some hand-specific characteristics, such as date, time, and the size of the mandatory bets (stake). For all four betting rounds, we further observe the actions of all active players, that is, players that have not folded their hands. On average, a player in our sample plays 2,159 hands and faces 53.48 two-person showdown situations during

the respective period. The average stack at the beginning of the hand is \$115.1, ranging from \$0.1 to \$81,644.⁸

For additional analyses and robustness checks, we extract several player-specific characteristics from the raw data, including the number of hands played over the observation period, the amount won/loss over the entire period (including hands with no showdown), and the average winning probability in showdown situations. As a result of the rake collected by Pokerstars, a player makes an average loss of \$0.088 per hand played. More details and the respective summary statistics are provided in Appendix A.4.

4 Results

We first describe our descriptive results, then our main regression analyses before we discuss in how far our results vary across different subgroups.

4.1 Descriptives


As we focus on situations where two players are in a showdown, that is, winning the pot is a complementary event for both opponents, each right-skewed lottery ($\pi < 0.5$) has exactly one left-skewed lottery ($1 - \pi > 0.5$) as a complement. The variance of the opponents' binary lotteries are identical, $V_1 = V_2 = \pi(1 - \pi)x^2$. Accordingly, the opponent of a player facing a binary risk with skewness S faces a lottery with a skewness of $-S$. Thus, we can directly analyze insurance choice frequencies for different signs of skewness, but with constant variance.

First, we group the observations conditional on the sign of skewness and find that individuals who face a negatively-skewed lottery choose the insurance option in 20.0% of cases. In contrast, individuals who face a positively-skewed risk do so in only 14.2% of the cases. This difference is statistically significant (p -value < 0.0001 , two-sided Welch's t -test, assuming independent samples) and in line with a preference for positive skewness.

Second, Figure 2 illustrates the choice frequencies of the insurance option for different ranges of ex-ante winning probabilities with a constant average variance. The horizontal axis depicts the winning probability range of right-skewed lotteries and the loss probability of the complementary left-skewed lotteries. For example, the first red bar at the left illustrates the insurance choice frequency for lotteries with a winning probability between 0 and 0.1, while the neighboring blue bar plots the frequency for lotteries with winning probabilities between 0.9 and 1, that is, loss probabilities between 0 and 0.1. In all subgroups, individuals who face a negatively (left-)skewed lottery choose the insurance option significantly more often than their opponents who face a positively (right-)skewed lottery with the same variance. The differences are sizable and statistically significant in all groups and range from 1.5% to 8.7% (see Table 11, Appendix A.4). The differences are smallest for ranges closer to $\pi = 0.5$, that is, for more symmetric lotteries. Insurance shares tend to be smaller for lotteries with a larger variance, suggesting a positive preference for variance by the average player. Interestingly, insurance shares seem

⁸The average stack and the average number of showdown situations are calculated using the 4,450,585 observations in our final data set. The number of hands played by a distinct player is based on all hands in our initial data set, including hands that did not result in a two-person showdown with a winning probability $\neq 0.5$.

Figure 2: Share of insurance choices for different winning probability ranges



Note: Figure 2 depicts the share of insurance choices depending on ex-ante winning probabilities. The probability space is divided into 10 equidistant segments. Right-skewed and complementary left-skewed risks with the same variance are grouped together (e.g., the first red bar at the left refers to the interval of right-skewed risks with winning probabilities in the range (0, 0.1] while the neighboring blue bar refers to the interval of left-skewed risks with loss probabilities in (0, 0.1]). For more details on observations and differences between groups see Table 11, Appendix A.4.

rather constant concerning the variance for right-skewed lotteries and tend to decrease for a higher variance for left-skewed lotteries.

While these results support a preference for skewness, they do have certain drawbacks, as they, for instance, do not account for different expected values of the gambles and other factors that may confound our results. We address these issues in our regression analyses.

4.2 Regression analyses

Empirical Strategy. We investigate how each moment of the underlying probability distribution influences individual insurance choices while holding the other moments constant. We follow Mitton and Vorkink (2007) in assuming that the different risk moments have a linear effect on utility. We do not have a clear prior regarding the influence of the expected value, given that both the safe option and the lottery exhibit roughly the same expected value. In contrast, we expect a positive (negative) sign for variance if individuals in our sample are, on average, risk-averse (risk-seeking). Skewness preferences imply a negative skewness coefficient, meaning that individuals choose the risky option more often for higher skewness. In our main specifications, we estimate the following reduced-form equation:

$$y_{i,j(t,z)} = \beta_0 + \beta_E E_j + \beta_V V_j + \beta_S S_j + \gamma \mathbf{Z}_i + \eta \mathbf{W}_j + \lambda_t + \psi_z + \epsilon_{i,j}, \quad (5)$$

where the dependent variable $y_{i,j(t,z)}$ is a binary indicator of whether player i chooses the insurance option in decision j . Each decision refers to a specific month t and a game with stake z . Variables E_j , V_j and S_j denote the expected value, variance, and skewness of the binary risk in decision j . λ_t are month fixed effects that control for month-specific factors constant across players, such as seasonality, adaptations over time, or COVID-19 effects. ψ_z captures stake fixed effects that account for the fact that a game with a higher stake directly implies higher average expected values and variance. If such fixed effects were not included, our coefficients would not only capture the effect of E_j and V_j on insurance choice but also heterogeneity between games with different stakes. Finally, $\epsilon_{i,j}$ denotes the error term.

To account for confounding factors related to the features of a particular hand, we control for hand-specific characteristics \mathbf{W}_j , including the amount of money the player started the hand with (that is, the stack), whether the player risked the entire stack during a particular showdown, the weekday, and the position of the respective player during a hand.⁹ The vector \mathbf{Z}_i includes a set of player-specific characteristics to control for different players' playing styles and experience levels. Player-specific characteristics are based on all poker hands in our data set (including those without a showdown) and cover the following variables: number of hands played, number of showdowns, profit or loss per 100 hands played, and the average winning probability in showdowns over all hands (summary statistics can be found in Appendix A.4).

While including player-specific controls addresses some endogeneity concerns, our estimated coefficients may still be biased if unobserved factors correlate with the type of lotteries individuals face. To address this issue, we exploit the panel structure of our data and include player-specific fixed effects α_i that control for all time-invariant heterogeneity across individuals. We extend the previous specification and estimate the following fixed effect regression:

$$y_{i,j(t,z)} = \beta_E E_j + \beta_V V_j + \beta_S S_j + \eta \mathbf{W}_j + \lambda_t + \psi_z + \alpha_i + \epsilon_{ij} \quad (6)$$

Regression results. Table 3 shows the estimated marginal effects from a linear probability model estimating Equations 5 and 6. To simplify the comparison of coefficients' magnitudes, we standardize the different moments in our main specifications.¹⁰ The signs and the p-values of the coefficients are largely unchanged if we use non-standardized variables (see Table 18, Appendix A.6). Similarly, our main results are unchanged if we estimate a Probit or a Logit model instead (see Table 19, Appendix A.6). Using a standardized dependent variable has the

⁹Two remarks on the hand-specific characteristics: i) in a two-person showdown, there is always one player that is all-in, thereby risking her entire stack, because otherwise, betting between the two players would still be possible, which rules out a showdown; ii) the position during a hand indicates when a player has to act during the hand, which may have important implications for the playing style and whether a player decides to play (i.e., voluntarily putting money in the pot) a particular set of hole cards or not. For example, players who already put money into the pot by posting the mandatory blinds tend to play a wider selection of hole cards, as the posted blinds count towards the necessary amount they have to call to see the first three community cards. Similarly, players who act last in each betting round (button) will have more information on opponents' actions when they act in future betting rounds, which usually increases the range of played sets of hole cards as well.

¹⁰We follow the literature and standardize the variables by computing the z-score, that is, we subtract the respective mean and scale the variable by the inverse of its standard deviation. This allows us to make a unit-independent comparison of the coefficient magnitudes.

additional advantage that the coefficient of the constant can be approximately interpreted as the average insurance choice shares in the particular (sub)sample.

Table 3: Regression results for full sample

	<i>Dependent variable:</i>			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Expected Value	0.002* (2.011)	0.003*** (3.358)	0.003*** (3.560)	0.007*** (7.166)
Variance	-0.0004 (-0.950)	-0.001 (-1.915)	-0.001 (-1.551)	-0.003*** (-4.757)
Skewness	-0.023*** (-17.117)	-0.023*** (-18.476)	-0.024*** (-18.890)	-0.023*** (-18.577)
Constant	0.171*** (41.932)	0.171*** (47.518)	0.171*** (47.792)	
Player-specific controls	No	Yes	Yes	No
Hand-specific controls	No	No	Yes	Yes
Player fixed effects	No	No	No	Yes
Observations	4,450,585	4,449,739	4,449,739	4,450,585

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5). The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. The number of observations in Columns 2 and 3 differs because the average winning probability (a player-specific control variable) is not available in 846 choice situations. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

If we only include the three moments as independent variables (see Column 1, Table 3), we find that a one standard deviation increase in skewness reduces the likelihood of choosing the insurance by 2.3 percentage points, or 13.4% compared to the average insurance choice likelihood (i.e., the mean dependent variable), while holding expected value and variance constant. The coefficient is statistically significant (p -value < 0.0001).

The impact of variance and expected value are significantly smaller. A one standard deviation rise in variance reduces insurance selection by an insignificant 0.04 percentage points. The negligible effect of the second moment is noteworthy in light of our substantial sample size, especially when contrasted with the skewness parameter's notable t-statistic of -17.117. A higher expected value increases the probability that the insurance is chosen by 0.2 percentage points. Note that this does not mean that Poker players dislike positive returns because the insurance value also increases with the risk's expected value. The positive coefficient indicates that players are more likely to choose the insurance option when facing lotteries with larger expected values.

Incorporating player-specific characteristics (Column 2) and hand-specific controls (Column 3) does not substantially alter the effects of expected value and skewness. The variance effect, while slightly larger, remains insignificant and substantially smaller than the coefficients

of the other two moments. The results of the player fixed effects regression model are presented in Column 4 of Table 3. The estimated effect of skewness does not change in magnitude compared to the base specification. The effects of expected value and variance increase, but variance's influence is still substantially smaller than skewness in absolute terms.

Overall, our regressions demonstrate a significant effect of skewness on insurance take-up. Despite extensive data, we fail to find a statistically detectable effect of variance—what is typically regarded as a risk's main property—in our base specifications. The estimated effect of skewness is negative and both statistically and economically significant. The absolute magnitude and the t-statistics of the standardized skewness coefficients are considerably larger than for the other two moments, suggesting a preeminent role of skewness preferences for individual risk-taking.

4.3 Heterogeneous effects

In this subsection, we estimate our main specifications for different potentially interesting subgroups. We split the full sample at the median of different player- and hand-specific characteristics, ensuring an even distribution of observations across the subsamples. Across all subsamples, we find evidence of skewness preferences, with the skewness coefficient being negative and significant in all specifications. For a list and summary statistics of the characteristics, see Tables 6–10 in Appendix A.4.

Players' experience. Players' experience is a potentially important dimension for understanding players decisions in showdown situation. Previous studies suggest that experienced poker players are more self-reflective, less affected by negative emotions, and make better decisions, by mathematical standards, than inexperienced players (Palomäki *et al.*, 2013, 2014). In our study, this suggests that the observed effects for experienced players are less likely due to a misperception of the underlying lotteries or due to impulsive choice but rather reflect a preference for different risk moments. Moreover, we have more repeated observations for more experienced players (for varying risks and payouts), which increases the power of our fixed effect regression model (see Table 12, Appendix A.5).

Experienced players face more than 421 showdowns. This threshold number balances the observations in the subsamples of inexperienced and experienced players. Experienced players opt for insurance in 22.0% of the cases when facing a left-skewed and 14.2% of the cases when facing a right-skewed lottery in the showdown. In contrast, the difference in insurance choice ratios between the left- and the right-skewed lotteries is considerably smaller for inexperienced players (17.9% vs. 14.2%). This pattern is consistent across different levels of lottery variance (Panels A and B of Figure 6, Appendix A.5). Again, for both subgroups, differences in insurance choices between left- and right-skewed lotteries narrows, as skewness approaches zero (see Table 13, Appendix A.5).

In our regression analyses, we also find large heterogeneity in skewness preferences between experienced and inexperienced players. For experienced players (as measured via the number of showdowns), increasing skewness by one standard deviation decreases the likelihood

Table 4: Regression results for different levels of player experience

<i>Dependent variable:</i>				
Insurance choice dummy				
	# showdowns		# hands	
	≤ 421	> 421	$\leq 13,861$	$> 13,861$
	(1)	(2)	(3)	(4)
Expected Value	0.006*** (6.393)	-0.002 (-1.882)	0.007*** (5.052)	-0.002 (-1.759)
Variance	-0.002*** (-3.580)	0.001 (1.925)	-0.003** (-3.234)	0.001* (2.121)
Skewness	-0.013*** (-18.021)	-0.033*** (-13.157)	-0.014*** (-18.389)	-0.032*** (-12.895)
Constant	0.160*** (113.620)	0.182*** (23.004)	0.164*** (110.186)	0.178*** (22.522)
Observations	2,228,808	2,221,777	2,225,785	2,224,800
Unique players	81,278	1,941	80,936	2,283

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5) for different subsamples. The full sample is split at the median of two different measures of player experience (as in Figure 6): i) the number of observed showdown situations per player (Columns 1 and 2); and ii) the total number of played hands by each player, including those without a showdown (Columns 3 and 4). At this place, we only present results from the specification without additional control variables or individual fixed effects (equivalent to Column 1 of Table 3). The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

of choosing the insurance option by 3.3 percentage points, which is equivalent to a decrease of 18.1% of the mean dependent variable in the respective subsample (Column 2, Table 4). For inexperienced players, the estimated effect of skewness on individuals' risk-taking is less than half in size (Column 1 Table 4). This difference persists when we include individual fixed effects and hand-specific control variables (see Table 12, Appendix A.5).

The coefficients of the other two moments also differ considerably for experienced and inexperienced players. The effect of variance is quite small in all subsamples and has different signs for experienced and inexperienced players in the basic specification, similar to the expected value coefficient. These results further underscore our earlier observation that, in our sample, skewness is a more influential factor than variance.

In Panels C and D of Figure 6 (Appendix A.5) and Columns 3 and 4 of Table 4, we split the sample according to an alternative measure of experience: the total number of hands played by an individual player (including those without a showdown). The cutoff value is equal to 13,861 played hands. The differences between the two subgroups align with our previous find-

ings and remain essentially unchanged if we include fixed effects (see Table 12, Appendix A.5). Moreover, the impact of other risk moments is similar across both measures of experience.

Players' success. Skill partly drives poker players' success, which separates poker from games of pure chance, such as roulette (Potter van Loon *et al.*, 2015; Duersch *et al.*, 2020). While skill does not affect the outcome of the binary lotteries we study in this paper, it affects the outcome of the preceding strategic interactions taking place under imperfect information. Skilled players typically achieve higher net returns due to better judgment of hand strength, more accurate predictions of opponents' hands, and adjusted betting strategies. Consequently, there are likely systematic differences in playing motives, risk attitudes, and preferences between successful and unsuccessful players.

We categorize players into successful and unsuccessful based on their aggregated net profits over our observation period, including results from non-showdown hands. Successful players are those with positive aggregated net profits over our observation period and unsuccessful players those with aggregated net losses. To account for the fact that profits can be the result of sheer luck if a player only plays a few hands, we also examine the heterogeneity in the subsample of experienced players only, who play at least 13,861 hands over our observation period. We distinguish between sophisticated players and recreational players. Sophisticated players are experienced successful players, while we define recreational players as experienced unsuccessful players. Like sophisticated individual investors, sophisticated poker players are more likely to reflect upon and adjust their strategy. Sophisticated players are less likely to play for recreational reasons only and are thus potentially more similar to financial experts. In contrast, recreational players and their underlying risk attitudes may be more comparable to speculative retail investors or people who gamble in a casino. It is noteworthy that successful players are often experienced. In our data set, the median experienced player makes an average loss of \$3.91 per hundred hands, while the median inexperienced player makes an average loss of \$20.43.

Table 14, Appendix A.5, presents the results of our subsample analysis. We find strong evidence of skewness preferences in both subsamples. Notably, unsuccessful players exhibit much stronger skewness preferences, particularly those in the recreational subgroup. In this group, a one standard deviation increase in skewness reduces the likelihood of choosing insurance by about five percentage points, consistent across both the base and fixed effects models (Columns 5 and 7). The effects of the other two moments align with those observed in the full sample. For successful and sophisticated players (Columns 2, 4, 6, and 8), the absolute effect of increased skewness on insurance choice is smaller, at only 1.0-1.1 percentage points, but remains significant across all models. In contrast, the expected value and variance effects are notably weaker and become statistically insignificant in the fixed effect model for sophisticated players (Column 8).

The subsamples also differ in average insurance rates. Successful players choose insurance in about 11.4% of cases, compared to 19.4% for unsuccessful players. The difference in insurance shares increases if we consider recreational and sophisticated players (23.2% vs. 10.6%). The differences, both in terms of insurance shares and risk preferences, are consistent with the idea that risk neutrality is advantageous for long-term earnings. Therefore, successful poker players

are expected to show less sensitivity to risk moments. Nonetheless, skewness still significantly influences the insurance decisions of successful players, indicating its importance even among those who are more risk-neutral.

Large monetary stakes. It is particularly interesting to see whether our results also hold when we restrict our analysis to large monetary stakes, that is, for decisions involving pots of more than \$100. The median lottery in this subsample has an expected value of \$95.45, with a mean of \$252.93. In this subsample, players selected the insurance option in 104,390 of 892,191 instances, yielding an insurance rate of 11.7%. Notably, players facing negatively-skewed lotteries chose the insurance in 13.7% of the cases, while those with positively-skewed risks did so in 9.7% of the cases.

Our regression results, which are presented in Table 15, Appendix A.5, show that with large monetary stakes, the absolute magnitude of the skewness coefficients moderately decrease to 1.6 to 1.9 percentage points, yet remain statistically significant across all specifications. Thus, our main conclusions remain unchanged when restricting our analysis to decisions with larger monetary stakes.

5 Limitations of our framework and robustness

In this section, we show the robustness of our results by addressing potential limitations and shortcomings of our setup. We first address the fact that, in our setting, lotteries are not always binary. Second, we disentangle skewness preferences from loss aversion, both of which could potentially drive our findings. Third, we address miscellaneous other issues.

5.1 Lotteries are not always binary

One limitation of our setup is that the underlying risks are, strictly speaking, not always binary, as split pots can occur. Split pots arise when players hold the same best five-card hand after all community cards are dealt. In this case, each involved player is awarded half of the pot. In our sample, 6.6% of all showdown situations result in a split pot. As players' best five-card hand includes precisely two of their private cards and three of the community cards, split pots are rare in Omaha Poker cash games compared to other Poker games. This is one reason why it is advantageous to focus on this particular poker variant in our study.

In scenarios where split pots are possible, players face the following trinary lottery: $L = (x, \pi; \frac{1}{2}x, \mu; 0, 1 - \mu - \pi)$, where μ is the ex-ante probability of a split pot and π is the ex-ante probability of winning the entire pot. In our data, μ and π are not independently observable, neither for the players nor for us. In fact, we only observe a "payout-weighted" winning probability $\tilde{\pi} = \frac{1}{2}\mu + \pi$, which can be understood as the percentage of the pot the player is expected to win. If a split pot is possible, the agent thus sees the binary lottery: $\tilde{L} = (x, \tilde{\pi}; 0, 1 - \tilde{\pi})$, for which $E(L) = E(\tilde{L})$. In Appendix A.3, we present an example of a choice situation where a split pot is possible and explain in detail how the payout-weighted probabilities are calculated. If no split pot is possible ex-ante ($\mu = 0$), which is true for the majority of hands, both lotter-

ies are equivalent ($L = \tilde{L}$). Moreover, in every showdown situation, for the player that faces the right-skewed and the one that faces the left-skewed risk, the probability of a split pot is the same, the variance of the lottery is the same, and the absolute value of the lottery's skewness is the same. Consequently, the possibility of split pots should not systematically confound our estimated skewness effects.

Notably, there are situations where the likelihood of a split pot is equal to one ($\mu = 1$), where the hand will result in a split pot irrespective of the remaining cards drawn from the deck. In these situations, the weighted probability we observe is equal to 0.5, but no insurance option was offered as the hand outcome is deterministic. Our dataset does not allow for a clear distinction between these deterministic scenarios and those with a weighted probability of 0.5. To address this issue, we exclude all observations with a weighted probability of 0.5. Most of these observations refer to a situation where the underlying lottery is degenerate, and no insurance option was offered.¹¹ Moreover, a winning probability of 0.5 implies that the underlying risk is not skewed but symmetric.

In response to the issue of split pots, we run our regression analyses for a subsample that excludes all observations that resulted in a split pot. The results are illustrated in Table 24, Appendix A.6. The estimated coefficients and p-values of all three moments remain nearly unchanged compared to our main regressions (Table 3). Note, however, that this approach only excludes observations that ultimately resulted in a split pot and not all showdowns where a split pot is possible ex-ante. However, the robustness of our estimated effects reassures us that our results are not confounded by the possibility of split pots.

5.2 Disentangling skewness preferences and loss aversion

Loss aversion, a key element of prospect theory (Kahneman and Tversky, 1979), posits that individuals assess payoffs relative to a specific reference point and have a preference for avoiding a loss over attaining a similarly sized gain, defined relative to this reference point. In our setting, the amount contributed by a player to the pot could serve as a natural reference point. According to loss aversion, players then avoid the insurance option when the insurance pays less than the player contributed to the pot, as this would represent a safe loss. In the showdown situations that we analyze, players facing left-skewed risks usually recoup their contribution with the insurance option and thus realize a gain.¹² Conversely, players facing right-skewed risks typically incur a loss when choosing the insurance option due to their low winning probability and thus may prefer to take the risk. However, there are exceptions due to "dead" money from players who folded their hands prior to the showdown. These instances of players choosing between a right-skewed risk and an insurance option that offers a gain relative to their pot contributions allow us to distinguish between loss aversion and skewness preferences.

¹¹This is implied by the following observation: In total, there are 111,507 observations (out of 4,562,092) with an observed probability of 0.5. 95,401 of those observations result in a split pot. In these 95,401 cases, the insurance option was chosen in 830 cases (0.9%). In the 16,106 cases that did not result in a split pot, the insurance was chosen in 3,528 cases (21.9%).

¹²Because of the rake that is deducted from the pot by Pokerstars, there are some situations where players facing a left-skewed lottery are, in fact, in the loss domain.

As a first step, we compare insurance choices between players for whom the insurance choice represents a gain and for whom it represents a loss. To achieve this, we limit the range of winning probabilities to between 0.3 and 0.5, ensuring a set of similar lotteries. There are only very few observations below a winning probability of 0.3 where players obtain a gain from choosing the insurance option. Above 0.5, there are nearly no observations where the insurance option implies a loss. Our qualitative results, however, do not change if we expand the probability range. Our analysis reveals significant differences in average insurance shares (see Table 16, Appendix A.5). Players whose insurance option represents a gain have an average insurance share of 18.2%, compared to 13.3% for players whose insurance implies a loss. This disparity in insurance choices is evident across different subsamples analyzed above and for pots larger than \$100. The magnitude of these differences—when compared to the average insurance share within each subsample—remains notably stable across the various groups.

Table 5: Regression results for full sample including the loss aversion dummy

	<i>Dependent variable:</i>			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Expected Value	−0.001 (−1.292)	0.0003 (0.274)	−0.00004 (−0.043)	0.006*** (6.026)
Variance	0.001* (2.048)	0.001 (1.129)	0.001 (1.221)	−0.002** (−3.121)
Skewness	−0.009*** (−10.766)	−0.009*** (−11.025)	−0.009*** (−11.158)	−0.010*** (−12.444)
Loss aversion dummy	−0.025*** (−18.881)	−0.025*** (−21.048)	−0.026*** (−21.694)	−0.023*** (−19.315)
Constant	0.171*** (41.926)	0.171*** (47.538)	0.171*** (47.805)	
Player-specific controls	No	Yes	Yes	No
Hand-specific controls	No	No	Yes	Yes
Player fixed effects	No	No	No	Yes
Observations	4,417,835	4,417,000	4,417,000	4,417,835

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5), including the loss aversion dummy. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses. Standard errors for the fixed effects regression are clustered at the individual level. Number of observations in Columns 2 and 3 differ because the average winning probability, one of the player-specific control variables, is not available in 835 choice situations. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Next, we conduct regression analyses to explore these patterns further. Our analyses employ the same specifications as above, additionally including a "loss aversion" dummy variable indicating whether the insurance option represents a loss relative to the pot contribution. The coefficient on this dummy serves as our metric for loss aversion, indicating the change in the

likelihood of selecting insurance when it represents a loss as opposed to a gain. Everything else equal, the baseline model using the full sample shows a 2.5 percentage point decrease in the insurance choice likelihood when the insurance represents a loss (Table 5). This effect remains stable across various specifications. Similar to the skewness effect, the effect of loss aversion is more pronounced for players who participated in more hands and unsuccessful players who experienced aggregated real money losses during our observation period (see Table 17, Appendix A.5). For pots exceeding \$100, the absolute size of the estimated effect of the loss aversion dummy slightly decreases from 2.5 to about 1.9 percentage points. Including the loss aversion dummy also affects our estimates of different risk moments. While reduced to approximately one percentage point, the skewness effect is still highly significant, unlike the effects of variance and expected value (Table 5). We again estimate the strongest skewness effects for the subsample of experienced and unsuccessful players. However, the skewness effects remain significant for inexperienced and successful players, and when considering only pots over \$100 (Table 17, Appendix A.5). Our findings suggest evidence of both skewness preferences and loss aversion.

5.3 Additional robustness checks

This subsection provides additional robustness checks to address other potential endogeneity issues and limitations. First, as mentioned above, we estimate Logit and Probit models to account for the inherent non-linear relationship between our binary outcome variable and our independent variables. The results are consistent with our main findings. The estimated average marginal effects of skewness, as well as their significance, remain nearly unchanged in both specifications (see Table 19, Appendix A.6).

Second, the first two moments of the lotteries depend on the net pot size. Our observations differ considerably in net pot sizes leading to substantial tails in the distribution of lotteries' expected value and variance. To make sure that this dispersion in pot sizes does not drive our results, we conduct our analyses using a normalized measure of lotteries' volatility, which is independent of the pot size: the "coefficient of variation" (CV) of the lotteries, which can be understood as the inverse of the "Sharpe ratio" of the lotteries.¹³ The CV is defined as the ratio of the standard deviation to the mean: $\sqrt{Var(L)}/E(L) = \sqrt{\pi(1-\pi)x^2}/\pi x = \sqrt{(1-p)/p}$. This measure is dimensionless and commonly used in finance and economics (similar to the Sharpe ratio, Sharpe, 1994) and in psychology (e.g., Weber *et al.*, 2004). We estimate the same regression equation as above with the only difference that we replace the expected value and the variance of the lottery with its CV. Table 20 in the Appendix shows the estimated marginal effects. The skewness coefficient slightly increases and remains highly statistically significant. The coefficient of the CV is slightly positive in all specifications and statistically significant at the 1%-level (except for the specification with player fixed effects). As an additional corroboration that our results are not driven by outliers, we run our main specification with samples trimmed at the 1%- and the 99%-percentiles of the lotteries' net pot. The results are shown in Table 21.

¹³Note that the Sharpe ratio in Finance is usually defined in terms of the difference between a risky investment's return and the risk-free return.

The skewness coefficient remains almost unchanged. The variance coefficient slightly increases in absolute terms to -0.5 percentage points and is now also highly significant. The coefficients are comparable if we trim the sample with respect to the expected at the 5%- and 95%-percentiles or if we winsorize the samples instead of trimming (results not shown).

Third, we run the same regressions for a sample that only includes players that face both types of lotteries—left- and right-skewed—at least once. This does also not change our results (see Table 22, Appendix A.6). Thus, our skewness effects are not driven by a fundamental difference between individuals facing left-skewed and individuals facing right-skewed risks.

Fourth, we estimate Equations 5 and 6, excluding all observations of players who never or always choose the insurance option. This rules out that our results are driven by the fact that players who always choose the insurance option face fundamentally different binary risks than players who never choose the insurance option. Limiting our sample to those players increases the estimated skewness effect (Table 23, Appendix A.6). Again, the (absolute) effect size and t-statistics of the skewness variable clearly exceed the estimates of the other moments in those specifications.

Finally, it is conceivable that the insurance decision is affected by other factors connected to the surrounding poker game. We provide additional robustness checks in the appendix controlling for two features of the surrounding poker game (Table 25, Appendix A.6). First, we control for the players' last action in a hand. A player's last action is either betting or calling, whereas there is always one caller and one bettor per hand. The caller knows the hand will end in a showdown, but the bettor is unsure whether their opponent will call, which creates uncertainty about reaching a showdown. This difference between players can affect their emotional state when deciding whether to choose the insurance. Second, we control for whether the opponent in the showdown took the insurance or not. While the opponent's insurance decision does not influence the outcome of the showdown, as noted above, players in a game observe the insurance decision of the other players, allowing for a potential behavioral response to the other's decision. Including these control variables in our regression analyses does not affect our main results. Moreover, the coefficients on these two additional control variables are rather small and insignificant in most specifications.

6 Discussion and Conclusion

The introduction of the insurance option in online poker allows us to cleanly test for skewness preferences in a large set of observational data among individuals who are rather experienced in choice under risk. We detect a strong and robust effect of skewness on risk-taking. Our results complement, for instance, recent survey findings (Holzmeister *et al.*, 2020) whereby skewness is the only moment that systematically affects financial professionals' perception of financial risk. We substantiate this finding in a real-world setting with a comprehensive data set of strongly incentivized choices.

Arguably, offering insurance (as reflected, in our case, by the all-in cashout) against negatively skewed risks as faced in a showdown situation increases the attractiveness of gambling. This way, players can enjoy positively skewed risks while being insured against negatively

skewed risks, which caters to players with skewness preferences. If such skewness preferences do not reflect true preferences, but are bias-driven as suggested, for instance, by the salience literature (Bordalo *et al.*, 2022; Dertwinkel-Kalt and Köster, forthcoming), the introduction of the all-in cashout induces players to gamble *excessively*. The online casino’s profits then increase due to the introduction of this option, but even though consumers demand it, their surplus potentially decreases.

As we analyze poker players, our study relies on a selective sample of people—a feature that is shared by most field studies on risky choice, which are restricted to, for instance, financial investors (e.g., Boyer *et al.*, 2010; Conrad *et al.*, 2013; Lin and Liu, 2018), game show participants (e.g., Gertner, 1993; Post *et al.*, 2008), bettors (e.g., Snowberg and Wolfers, 2010; Andrikogiannopoulou and Papakonstantinou, 2020), or people buying auto insurance (e.g., Cohen and Einav, 2007). In our case, the selected sample has advantages and disadvantages. Online poker may attract individuals with non-representative (risk) preferences, which is backed by the observation that the overall insurance take-up is rather low and playing online poker has, on average, a negative expected return due to the fee taken by the platform providers (on average a player loses around \$0.088 per hand). Past studies suggest that online poker players are more likely to be younger and predominantly male compared to the general population (Barrault and Varescon, 2016). While not necessarily applicable to the general population, our results could be rather informative for individuals that self-select into risky choices in other instances, such as individual investors—particularly those with a large propensity to invest in lottery-type stocks (Kumar, 2009; Han and Kumar, 2013) or cryptocurrencies (Hackethal *et al.*, 2022)—, bettors (e.g., Andrikogiannopoulou and Papakonstantinou, 2020; Moskowitz, 2021), and entrepreneurs (Moskowitz and Vissing-Jørgensen, 2002).

More generally, our results suggest that payoff skewness or lottery-like features are important drivers for risk-taking and (asset) prices, particularly in markets that are predominately populated by young speculative individuals, such as the markets for crypto assets. In addition, our findings have important real-world implications beyond asset pricing and risk-taking, as skewness preferences affect career choices and may explain recent phenomena, such as the boom of tech start-ups or the popularity of extended warranties for many durable goods, even when generous base warranties are in place (Lee and Venkataraman, 2022).

References

- ALLAIS, M. (1953). Le comportement de l’homme rationnel devant le risque: critique des postulats et axiomes de l’école américaine. *Econometrica*, **21**, 503–546.
- ANDRIKOGIANNOPOULOU, A. and PAPA KONSTANTINO U, F. (2020). History-dependent risk preferences: Evidence from individual choices and implications for the disposition effect. *Review of Financial Studies*, **33** (8), 3674–3718.
- BALI, T. G., CAKICI, N. and WHITELAW, R. F. (2011). Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, **99** (2), 427–446.

- BARBERIS, N. (2012). A model of casino gambling. *Management Science*, **58** (1), 35–51.
- BARRAULT, S. and VARESCON, I. (2016). Online and live regular poker players: Do they differ in impulsive sensation seeking and gambling practice? *Journal of Behavioral Addictions*, **5** (1), 41–50.
- BARSEGHYAN, L., MOLINARI, F., O'DONOGHUE, T. and TEITELBAUM, J. C. (2013). The nature of risk preferences: Evidence from insurance choices. *American Economic Review*, **103** (6), 2499–2529.
- BELL, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, **30** (5), 961–981.
- BERKHOUT, P., HARTOG, J. and WEBBINK, D. (2010). Compensation for earnings risk under worker heterogeneity. *Southern Economic Journal*, **76** (3), 762–790.
- BORDALO, P., GENNAIOLI, N. and SHLEIFER, A. (2022). Saliency. *Annual Review of Economics*, **14**, 521–544.
- , NICOLA GENNAIOLI and SHLEIFER, A. (2012). Saliency theory of choice under risk. *Quarterly Journal of Economics*, **127** (3), 1243–1285.
- , — and SHLEIFER, A. (2013). Saliency and asset prices. *American Economic Review, Papers & Proceedings*, **103** (3), 623–628.
- BOYER, B., MITTON, T. and VORKINK, K. (2010). Expected idiosyncratic skewness. *Review of Financial Studies*, **23** (1), 169–202.
- CHUNHACHINDA, P., DANDAPANI, K., HAMID, S. and PRAKASH, A. J. (1997). Portfolio selection and skewness: Evidence from international stock markets. *Journal of Banking & Finance*, **21** (2), 143–167.
- COHEN, A. and EINAV, L. (2007). Estimating risk preferences from deductible choice. *American Economic Review*, **97** (3), 745–788.
- CONRAD, J., DITTMAR, R. F. and GHYSELS, E. (2013). Ex ante skewness and expected stock returns. *Journal of Finance*, **68** (1), 85–124.
- DEWINKEL-KALT, M., EBERT, S. and KÖSTER, M. (2023). On correlated lotteries in economic applications. *Journal of Economic Behavior & Organization*, **215**, 292–306.
- and KÖSTER, M. (2020). Saliency and skewness preferences. *Journal of the European Economic Association*, **18** (5), 2057–2107.
- and KÖSTER, M. (forthcoming). Salient cues and complexity. *Management Science*.
- DUERSCH, P., LAMBRECHT, M. and OECHSSLER, J. (2020). Measuring skill and chance in games. *European Economic Review*, **127**, 103472.
- EBERT, S. (2015). On skewed risks in economic models and experiments. *Journal of Economic Behavior & Organization*, **112**, 85–97.

- and KAREHNKE, P. (2020). Skewness preferences in choice under risk. *Available at SSRN*.
- and WIESEN, D. (2011). Testing for prudence and skewness seeking. *Management Science*, **57** (7), 1334–1349.
- EIL, D. and LIEN, J. W. (2014). Staying ahead and getting even: Risk attitudes of experienced poker players. *Games and Economic Behavior*, **87**, 50–69.
- FAMA, E. F. and FRENCH, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance*, **47** (2), 427–465.
- FIEDLER, I. C. and ROCK, J.-P. (2009). Quantifying skill in games—theory and empirical evidence for poker. *Gaming Law Review and Economics*, **13** (1), 50–57.
- GARRETT, T. A. and SOBEL, R. S. (1999). Gamblers favor skewness, not risk: Further evidence from united states' lottery games. *Economics Letters*, **63** (1), 85–90.
- GERTNER, R. (1993). Game shows and economic behavior: risk-taking on "card sharks". *Quarterly Journal of Economics*, **108** (2), 507–521.
- GOLEC, J. and TAMARKIN, M. (1998). Bettors love skewness, not risk, at the horse track. *Journal of Political Economy*, **106** (1), 205–225.
- GROVE, W. A., JETTER, M. and PAPPS, K. L. (2021). Career lotto? Labor supply in a superstar market. *Journal of Economic Behavior & Organization*, **183**, 362–376.
- GUL, F. (1991). A theory of disappointment aversion. *Econometrica*, **59**, 667–686.
- HACKETHAL, A., HANSPAL, T., LAMMER, D. M. and RINK, K. (2022). The characteristics and portfolio behavior of bitcoin investors: Evidence from indirect cryptocurrency investments. *Review of Finance*, **26**, 855–898.
- HAN, B. and KUMAR, A. (2013). Speculative retail trading and asset prices. *Journal of Financial and Quantitative Analysis*, **48** (2), 377–404.
- HARTOG, J. and VIJVERBERG, W. P. (2007). On compensation for risk aversion and skewness affection in wages. *Labour Economics*, **14** (6), 938–956.
- HOLZMEISTER, F., HUBER, J., KIRCHLER, M., LINDNER, F., WEITZEL, U. and ZEISBERGER, S. (2020). What drives risk perception? A global survey with financial professionals and laypeople. *Management Science*, **66** (9), 3977–4002.
- INMAN, J. J., DYER, J. S. and JIA, J. (1997). A generalized utility model of disappointment and regret effects on post-choice valuation. *Marketing Science*, **16** (2), 97–111.
- JONDEAU, E., ZHANG, Q. and ZHU, X. (2019). Average skewness matters. *Journal of Financial Economics*, **134** (1), 29–47.
- KAHNEMAN, D. and TVERSKY, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47** (2), 263–291.

- KUMAR, A. (2009). Who gambles in the stock market? *Journal of Finance*, **64** (4), 1889–1933.
- LEE, H.-T. and VENKATARAMAN, S. (2022). A reference-dependent-preferences-based model of extended warranty purchase decisions. *Journal of Marketing Research*, **59** (3), 641–658.
- LIN, T.-C. and LIU, X. (2018). Skewness, individual investor preference, and the cross-section of stock returns. *Review of Finance*, **22** (5), 1841–1876.
- LOOMES, G. and SUGDEN, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, **92** (368), 805–824.
- and — (1987a). Some implications of a more general form of regret theory. *Journal of Economic Theory*, **41** (2), 270–287.
- and — (1987b). Testing for regret and disappointment in choice under uncertainty. *Economic Journal*, **97**, 118–129.
- MARKOWITZ, H. (1952). The utility of wealth. *Journal of Political Economy*, **60** (2), 151–158.
- MITTON, T. and VORKINK, K. (2007). Equilibrium underdiversification and the preference for skewness. *Review of Financial Studies*, **20** (4), 1255–1288.
- MOSKOWITZ, T. J. (2021). Asset pricing and sports betting. *Journal of Finance*, **76** (6), 3153–3209.
- and VISSING-JØRGENSEN, A. (2002). The returns to entrepreneurial investment: A private equity premium puzzle? *American Economic Review*, **92** (4), 745–778.
- PALOMÄKI, J., LAAKASUO, M. and SALMELA, M. (2013). “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.
- , — and SALMELA, M. (2014). Losing more by losing it: Poker experience, sensitivity to losses and tilting severity. *Journal of Gambling Studies*, **30** (1), 187–200.
- POST, T., VAN DEN ASSEM, M. J., BALTUSSEN, G. and THALER, R. H. (2008). Deal or no deal? decision making under risk in a large-payoff game show. *American Economic Review*, **98** (1), 38–71.
- POTTER VAN LOON, R. J., VAN DEN ASSEM, M. J. and VAN DOLDER, D. (2015). Beyond chance? the persistence of performance in online poker. *PLOS one*, **10** (3), e0115479.
- PRAKASH, A. J., CHANG, C.-H. and PACTWA, T. E. (2003). Selecting a portfolio with skewness: Recent evidence from us, european, and latin american equity markets. *Journal of Banking & Finance*, **27** (7), 1375–1390.
- PRIMEDOPE (2023). Largest poker sites and online poker traffic history. <https://www.primedope.com/largest-poker-sites/>. Last accessed January 13, 2024.
- SHARPE, W. F. (1994). The Sharpe ratio. *Journal of Portfolio Management*, **21**, 49–59.

- SMITH, G., LEVERE, M. and KURTZMAN, R. (2009). Poker player behavior after big wins and big losses. *Management Science*, **55** (9), 1547–1555.
- SNOWBERG, E. and WOLFERS, J. (2010). Explaining the favorite–long shot bias: Is it risk-love or misperceptions? *Journal of Political Economy*, **118** (4), 723–746.
- SYDNOR, J. (2010). (Over)insuring modest risks. *American Economic Journal: Applied Economics*, **2** (4), 177–199.
- THALER, R. H. and JOHNSON, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science*, **36** (6), 643–660.
- TRAUTMANN, S. T. and VAN DE KUILEN, G. (2018). Higher order risk attitudes: A review of experimental evidence. *European Economic Review*, **103**, 108–124.
- WEBER, E. U., SHAFIR, S. and BLAIS, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: risk as variance or coefficient of variation. *Psychological Review*, **111** (2), 430–445.

A Appendix

A.1 Poker hands ranking

The player with the highest ranked five-card hand, consisting of two private cards and three community cards, wins the pot in Omaha Poker cash games. The poker hand ranking is as follows (Source: Pokerstars: <https://www.pokerstars.eu/poker/games/rules/hand-rankings/>):

1. Straight Flush: Five cards in numerical order, all of identical suits. In the event of a tie, the highest rank at the top of the sequence wins. The best possible straight flush is known as a royal flush, which consists of the ace, king, queen, jack, and ten of a suit. A royal flush is an unbeatable hand.

2. Four of a Kind: Four cards of the same rank, and one side card or 'kicker.' In the event of a tie, the highest four of a kind wins. In community card games, where players have the same four of a kind, the highest fifth side card ('kicker') wins.

3. Full House: Three cards of the same rank, and two cards of a different, matching rank. In the event of a tie, the highest three matching cards wins the pot. In community card games, where players have the same three matching cards, the highest value of the two matching cards wins.

4. Flush: Five cards of the same suit. In the event of a tie, the player holding the highest ranked card wins. If necessary, the second-highest, third-highest, fourth-highest, and fifth-highest cards can be used to break the tie. If all five cards are the same rank, the pot is split. The suit itself is never used to break a tie in poker.

5. Straight: Five cards in sequence. In the event of a tie, the highest ranking card at the top of the sequence wins. Note: The Ace may be used at the top or bottom of the sequence, and is the only card which can act in this manner. A,K,Q,J,T is the highest (Ace high) straight; 5,4,3,2,A is the lowest (Five high) straight.

6. Three of a kind: Three cards of the same rank, and two unrelated side cards. In the event of a tie, the highest ranking three of a kind wins. In community card games, where players have the same three of a kind, the highest side card, and if necessary, the second-highest side card wins.

7. Two pair: Two cards of a matching rank, another two cards of a different matching rank, and one side card. In the event of a tie: Highest pair wins. If players have the same highest pair, highest second pair wins. If both players have two identical pairs, highest side card wins.

8. One pair: Two cards of a matching rank, and three unrelated side cards. In the event of a tie, the highest pair wins. If players have the same pair, the highest side card wins, and if necessary, the second-highest and third-highest side card can be used to break the tie.

9. High card: Any hand that does not qualify under a category listed above. In the event of a tie, the highest card wins, and if necessary, the second-highest, third-highest, fourth-highest, and smallest card can be used to break the tie.

A.2 Choice situation for a showdown before the flop

Figure 3 shows another example of a showdown situation, in particular of a constellation before any community cards have been dealt. The difference from the example in the main text is the different stage of the game when the showdown situation has occurred. The situation environment for the player is equivalent, that is, payouts and probabilities are clearly displayed on the player's screen and the decision the player takes is the same. Note that the players have the insurance option only once, namely in the moment of showdown.

Figure 3: Example of an all-in cashout situation before any community cards have been dealt (German software)



A.3 Showdowns with split pot possibility

Figure 4 shows an example of a showdown situation with one card to come. The difference from the example in the main text is that there is a split pot possibility. After the Turn, Player 1 holds the best five-card hand with a "Straight" (5-6-7-8-9). Player 4's best possible five-card hand is 8-8-8-K-9, three of a kind. Again, there are still 40 cards in the deck. Player 4 would win the entire pot if the board pairs, that is, if a King, 9, 5 or 8 is drawn, giving her a winning Full House. As Player 1 holds one 8 and one 5, there are seven cards in the remaining deck that would give Player 4 the winning hand. Accordingly, the likelihood for Player 4 to win the entire pot is $\pi = \frac{7}{40} = 0.175$. However, as player 4 also holds a 6 (and 8/9) in his hand, a 7 on the River would give her the same straight (5-6-7-8-9) as Player 1, which would result in a split pot. There are three 7s still in the deck, implying a probability for a split pot of: $\mu = \frac{3}{40} = 0.075$. As Player 4 would only win half of the pot in this case, $\frac{1}{2}\mu$ is added to the winning probability to get the "payout-weighted probability," or expected winning share of pot, that is displayed on

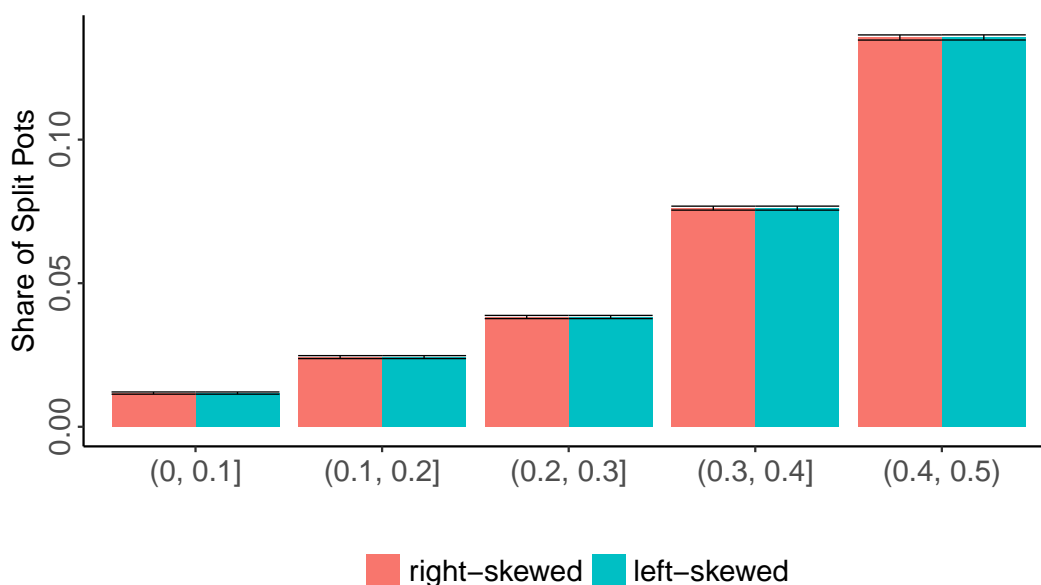
the screen: $\pi + \frac{1}{2}\mu = 0.175 + \frac{1}{2}0.075 = 0.2125$. Apart from that, the decision environment for the player is equivalent, that is, payouts are clearly displayed on the player's screen and the players only have the insurance option once, namely in the moment of showdown.

Figure 5 illustrates shares of hands that result in a split pot depending on expected winning shares of the pot.

Figure 4: Example of an all-in cashout situation on the Turn with a split pot possibility



Figure 5: Share of split pots depending on the expected winning share of the pot



Note: Figure 5 presents the share of hands that resulted in a split pot depending on expected winning shares of the pot. The expected winning share is equivalent to the "payout-weighted" winning probability introduced in Section 5.1. The shares are divided into 10 equidistant segments and right-skewed and complementary left-skewed risks are grouped together (e.g., the first red bar on the left refers to the interval of right-skewed risks with expected winning shares in the range (0,0.1] and the neighboring blue bar refers to the interval of left-skewed risks with expected winning shares in [0.9,1)).

A.4 Additional summary statistics and descriptives

Table 6: Summary statistics of player-specific characteristics

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Number of hands played	83,219	2,159.42	12,932.65	1.00	53.00	186.00	785.00	602,706.00
Number of experienced showdown situations	83,219	53.48	252.60	1.00	2.00	7.00	26.00	13,930.00
Average winning probability	82,470	0.44	0.15	0.00	0.37	0.45	0.52	1.00
Profit per hundred hands	83,219	-98.01	1,048.16	-145,158.40	-66.45	-19.73	-2.27	36,620.00

Note: The table reports summary statistics of all player-specific characteristics that we use in our empirical analysis. Characteristics are used, both, as control variables in our regressions analyses and to split the full sample for our subsample analyses (Section 5.3). The variable *Number of experienced showdown situations* is calculated using the 4,485,585 observations of showdown situations in our full sample. The other characteristics are based on all hands in our initial data set, which also includes hands that did not result in a two-person showdown with a winning probability of $\neq 0.5$. The statistics are calculated with equal weights on all players. For the profit per hundred hands variable this implies that players who have played fewer hands are heavily overweighted in the calculation of the summary statistics. As these players usually make high losses in few hands and then stop playing, we get a large discrepancy between the average profits per hundred hand across all hands and players, which is equal to $-\$8.88$, and the profit per hundred hands when the mean is calculated with equal weights on single players (as illustrated in the table). The number of observations (N) for average winning probability differs compared to other characteristics as these values are not available in the data for 749 players.

Tables 7-9 report summary statistics of all hand-specific characteristics that we use in our empirical analysis.

Table 7: Summary statistics of stakes & stacks

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Stake	4,450,585	1.195	5.505	0.100	0.100	0.250	0.500	100.000
Stack	4,450,585	115.088	628.627	0.100	10.390	25.240	63.960	81,643.93

Note: The table reports summary statistics of the stake (mandatory bets) and the stacks (money of each player at the beginning of the hand) in our sample. Values are measured in US-\$ terms.

Table 8: Summary statistics of the risk-all-stack dummy, the dummy indicating insurance choice of the other player and the dummy indicating whether the last action of a player was a call

Statistic	N	# Dummy = 1	# Dummy = 0	Mean	St. Dev.
Risk-All-Stack dummy	4,450,585	2,225,425	2,225,160	0.500	0.500
Other player's insurance dummy	4,450,585	761,577	3,689,008	0.171	0.377
Call dummy	4,450,585	2,171,584	2,279,001	0.488	0.500

Note: The table reports the number of hand situations where the respective player risks her entire stack (Risk-all-stack=1), where the other player has chosen the insurance and where the last action of the respective player was a call. As we do not directly observe whether a player risks her entire stack in a showdown, we approximate the risk-all-stack dummy with the player's stack, final pot and expected winning shares. Due to rounding and presence of mandatory bets of other players who are not involved in the showdown, there might be some individuals that wrongly end up in the subsample of individuals that do not risk their entire stack. The error margin should be small and should not confound the results.

Table 9: Frequencies of positions in a showdown situation

	BB	BTN	CO	EP	MP	SB
Frequency	915,826	902,579	773,689	362,749	632,235	863,507

Note: The table reports absolute frequencies of the different positions of the player in a showdown situation. These are namely: BB ("Big Blind"; person that has to post the big blind), BTN ("Button"; person that acts last in every betting round after the Flop), CO ("Cut-Off"; person that acts second last in every betting round after the Flop), SB ("Small Blind", person that has to post the big blind) as well as EP ("Early Position") and MP ("Middle Position") that refers to positions between the "Big Blind" and the "Cut-Off".

Table 10: Frequencies of weekdays

	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Frequency	610,961	610,300	617,142	616,895	641,434	682,997	670,856

Note: The table reports absolute frequencies of the weekdays when showdown situations have occurred.


Table 11: Differences in shares of insurance choice among right- and left-skewed risks for different ranges of winning probabilities

right-skewed interval (r)	left-skewed interval (l)	Obs.(r)	Obs.(l)	Δ in shares	t -statistic
(0.0, 0.1]	[0.9, 1.0)	298,984	298,984	0.087***	86.165
(0.1, 0.2]	[0.8, 0.9)	344,012	344,013	0.085***	88.607
(0.2, 0.3]	[0.7, 0.8)	462,970	463,002	0.076***	95.299
(0.3, 0.4]	[0.6, 0.7)	563,305	563,297	0.054***	78.788
(0.4, 0.5)	(0.5, 0.6)	555,986	556,032	0.015***	21.340

Note: The table reports the number of observations in each of the ten equidistant probability ranges illustrated in Figure 2. Column 5 reports the difference in the shares of insurance choice between complementary groups of left- and right-skewed risks. The imbalance in observations between the left- and right skewed intervals is due to rounding differences in winning probabilities. Column 6 reports t -statistics of the two-sided Welch's t -test, assuming independent samples.
*: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

A.5 Additional subsample analyses

Figure 6: Share of insurance choices depending on players' experience



Note: The figure depicts the share of insurance choices depending on ex-ante winning probabilities, equivalent to Figure 2, for different subsamples of relatively inexperienced players (Panels A and C) and subsamples of relatively experienced players (Panels B and D). The sample is split at the median of the respective experience measure. In Panels A and B, players' experience is measured by the number of observed showdowns in the sample. In contrast, Panels C and D use the total number of played hands by each player, including those without a showdown, as a measure of the player's experience. The probability space is divided into 10 equidistant segments. Right-skewed and complementary left-skewed risks with the same variance are grouped together. For more details on observations and differences between groups see Table 13, Appendix A.5.

Table 12: Regression results in subsamples, with fixed effects

<i>Dependent variable:</i>				
Insurance choice dummy				
	# showdowns		# hands	
	≤ 421	> 421	$\leq 13,861$	$> 13,861$
	(1)	(2)	(3)	(4)
Expected Value	0.011*** (7.758)	0.003** (2.604)	0.013*** (7.275)	0.002* (2.337)
Variance	-0.004*** (-4.947)	-0.001 (-1.912)	-0.005*** (-4.824)	-0.001 (-1.558)
Skewness	-0.012*** (-17.295)	-0.035*** (-14.656)	-0.013*** (-17.081)	-0.034*** (-14.180)
Player-specific controls	No	No	No	No
Hand-specific controls	Yes	Yes	Yes	Yes
Player fixed effects	Yes	Yes	Yes	Yes
Observations	2,228,808	2,221,777	2,225,785	2,224,800
Unique players	81,278	1,941	80,936	2,283

Note: The table reports OLS regression coefficients of our fixed effects specification (Equation 6) for different subsamples. Compared to Table 4, we include individual fixed effects and hand-specific characteristics as control variables (equivalent to Column 4 of Table 3). The full sample is split at the median of two different measures of player experience: i) the number of observed showdown situations in our sample per player (Columns 1 and 2); and ii) the total number of played hands by each player, including those without a showdown (Columns 3 and 4). The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Table 13: Differences in shares of insurance choice among right- and left-skewed risks

r	l	# showdowns		# hands	
		≤ 421	> 421	$\leq 13,861$	$> 13,861$
(1)	(2)	(3)	(4)	(5)	(6)
(0.0, 0.1]	[0.9, 1.0)	0.044*** (32.088)	0.132*** (89.251)	0.049*** (35.267)	0.128*** (87.543)
(0.1, 0.2]	[0.8, 0.9)	0.049*** (37.326)	0.121*** (86.976)	0.054*** (40.775)	0.117*** (84.803)
(0.2, 0.3]	[0.7, 0.8)	0.048*** (43.805)	0.103*** (89.307)	0.053*** (47.609)	0.098*** (86.359)
(0.3, 0.4]	[0.6, 0.7)	0.041*** (42.663)	0.065*** (66.679)	0.045*** (46.300)	0.061*** (63.359)
(0.4, 0.5)	(0.5, 0.6)	0.012*** (12.590)	0.016*** (16.507)	0.013** (13.815)	0.015*** (15.337)

Note: The table presents the difference in the shares of insurance choice between left- and right-skewed risks for different subsamples and ranges of ex-ante winning probabilities, as defined in Columns 1 and 2 and illustrated in Figure 6. The full sample is split at the median of two different measures of player experience: i) the number of observed showdown situations in our sample per player (Columns 3 and 4); and ii) the total number of played hands by each player, including those without a showdown (Columns 5 and 6). Corresponding t-statistics of the two-sided Welch's t -test, assuming independent samples, are provided in parentheses. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Table 14: Regression results in subsamples, depending on players' success

<i>Dependent variable:</i>								
Insurance choice dummy								
	All players				Experienced players only			
	without fixed effects		with fixed effects		without fixed effects		with fixed effects	
	Profit per hundred hands		Profit per hundred hands		Profit per hundred hands		Profit per hundred hands	
	≤ 0	> 0	≤ 0	> 0	≤ 0	> 0	≤ 0	> 0
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Expected Value	0.006*** (4.093)	0.002* (2.010)	0.011*** (5.599)	0.003*** (3.343)	0.002 (1.133)	0.002 (1.147)	0.007** (2.690)	0.002 (1.636)
Variance	-0.002* (-2.572)	-0.001 (-1.855)	-0.004*** (-3.512)	-0.002** (-3.088)	-0.0001 (-0.204)	-0.001 (-1.126)	-0.002 (-1.887)	-0.001 (-1.690)
Skewness	-0.029*** (-16.785)	-0.011*** (-7.057)	-0.027*** (-17.422)	-0.011*** (-7.172)	-0.049*** (-13.116)	-0.010*** (-4.936)	-0.048*** (-13.868)	-0.011*** (-5.264)
Constant	0.193*** (42.059)	0.114*** (15.142)			0.231*** (21.618)	0.105*** (10.482)		
Observations	3,192,498	1,258,087	3,192,498	1,258,087	1,289,004	935,796	1,289,004	935,796
Player-specific controls	No	No	No	No	No	No	No	No
Hand-specific controls	No	No	Yes	Yes	No	No	Yes	Yes
Player fixed effects	No	No	Yes	Yes	No	No	Yes	Yes
Observations	3,192,498	1,258,087	3,192,498	1,258,087	1,289,004	935,796	1,289,004	935,796
Unique players	65,886	17,333	65,886	17,333	1,472	811	1,472	811

Note: The table reports regression coefficients for basic OLS specification (Equation 5) in Columns 1, 2, 5, and 6 and for the setup including player fixed effects (Equation 6) in Columns 3, 4, 7 and 8. The baseline samples are split into unsuccessful (columns with an uneven number) and winning (columns with an even number) players. Columns 1 to 4 employ the full sample, and Columns 5 to 8 use the sample of experienced players as the baseline sample. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Table 15: Regression results for the subsample of hands with pot size > 100 US-\$

	<i>Dependent variable:</i>			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Expected Value	0.001 (0.719)	0.002 (1.415)	0.002 (1.56)	0.005*** (4.469)
Variance	0.0004 (0.617)	0.0001 (0.186)	0.0002 (0.356)	-0.001* (-2.272)
Skewness	-0.016*** (-10.783)	-0.019*** (-14.283)	-0.019*** (-14.391)	-0.018*** (-13.490)
Constant	0.117*** (24.866)	0.117*** (27.331)	0.117*** (27.462)	
Player-specific controls	No	Yes	Yes	No
Hand-specific controls	No	No	Yes	Yes
Player fixed effects	No	No	No	Yes
Observations	892,191	892,130	892,130	892,191

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5) for the subsample of hands with pot size > 100 US-\$. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. Number of observations in Columns 2 and 3 differ because the average winning probability, one of the player-specific control variables, is not available in 61 choice situations. *: p<0.1; **: p<0.05; ***: p<0.01.

Table 16: Insurances choice shares for gain and loss implying lotteries

Sample	0.99E higher than contributions	0.99E lower than contributions	Difference
	(1)	(2)	(3)
(A) Full Sample	0.182 (0.386) 88,454	0.133 (0.339) 1,021,519	-0.049 (-36.891)
(B) Experienced Players (# hands played > 13,861)	0.192 (0.394) 41,861	0.137 (0.344) 506,370	-0.055 (-27.510)
(C) Inexperienced Players (# hands played ≤ 13,861)	0.173 (0.379) 46,593	0.128 (0.334) 515,149	-0.045 (-24.914)
(D) Unsuccessful Players (Profit per 100 hands ≤ 0)	0.202 (0.402) 66,082	0.147 (0.354) 728,183	-0.055 (-34.137)
(E) Successful Players (Profit per 100 hands > 0)	0.122 (0.327) 22,372	0.096 (0.295) 293,336	-0.026 (-11.316)
(F) Large Pot Size (Pot Size > 100 US-\$)	0.140 (0.346) 14,946	0.091 (0.287) 221,731	-0.049 (-16.757)

Note: The table reports the difference in insurance choice shares between decisions where the offered insurance payout was larger than the players' pot contributions and decisions where the offered insurance payout was smaller than the contributions. The comparison is made for various subsamples. The winning probabilities in all subsamples are limited to the range (0.3, 0.5). Columns 1 and 2 report the respective insurance choice shares, followed by standard deviations in parentheses, and the number of observations for the respective sample. Column 3 reports the difference between the insurance choice shares in Columns 1 and 2, along with the corresponding t-statistic in parentheses. *: p<0.1; **: p<0.05; ***: p<0.01.

Table 17: Regression results including loss aversion dummy for different subsamples

<i>Dependent variable:</i>					
Insurance choice dummy					
	Experienced Players	Inexperienced Players	Unsuccessful Players	Successful Players	Large Pot Size
	(1)	(2)	(3)	(4)	(5)
Expected Value	-0.006*** (-4.512)	0.007*** (4.443)	0.004* (2.530)	0.001 (0.630)	-0.004* (-2.493)
Variance	0.003* (2.346)	-0.003* (-2.201)	-0.001 (-1.266)	-0.0003 (-0.400)	0.003** (2.756)
Skewness	-0.015*** (-9.547)	-0.003*** (-6.142)	-0.010*** (-10.161)	-0.005*** (-4.121)	-0.005*** (-5.367)
Loss aversion dummy	-0.029*** (-12.633)	-0.019*** (-24.364)	-0.022*** (-12.688)	-0.025*** (-16.013)	-0.019*** (-13.108)
Constant	0.178*** (22.527)	0.164*** (110.303)	0.193*** (42.18)	0.114*** (15.114)	0.117*** (24.768)
Player-specific controls	No	No	No	No	No
Hand-specific controls	No	No	No	No	No
Player fixed effects	No	No	No	No	No
Observations	2,207,911	2,209,924	3,170,340	1,247,495	877,118
Unique Players	2,283	80,936	65,886	17,333	28,090

Note: The table reports OLS regression coefficients for our main empirical specification (Equation 5), including an additional control for the loss aversion dummy, which indicates whether a player is in the winning or loss domain in the respective hand, across different subsamples. Columns 1 and 2 represent subsamples which reflect players' experience by splitting observations at the median of the total number of hands played for each player, Columns 3 and 4 represent subsamples which reflect players' success by splitting observations based on whether a player has achieved a positive profit per hundred games played or not and Column 5 represents the subsample which reflects hands with large pot size, specifically those exceeding 100 US-\$. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

A.6 Robustness checks

Table 18: Regression results for full sample, non-standardized variables

	<i>Dependent variable:</i>			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Expected Value	0.000005* (2.011)	0.000008*** (3.358)	0.000009*** (3.560)	0.00002*** (7.166)
Variance	-0.0000000002 (-0.950)	-0.0000000003 (-1.915)	-0.0000000003 (-1.551)	-0.000000001*** (-4.757)
Skewness	-0.010*** (-17.117)	-0.010*** (-18.476)	-0.011*** (-18.890)	-0.010*** (-18.577)
Constant	0.186*** (21.905)	0.064*** (3.379)	0.044* (2.228)	
Player-specific controls	No	Yes	Yes	No
Hand-specific controls	No	No	Yes	Yes
Player fixed effects	No	No	No	Yes
Observations	4,450,585	4,449,739	4,449,739	4,450,585

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5). Compared to Table 3, the independent variable enters the regression as non-standardized absolute values. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. Number of observations in Columns 2 and 3 differ because the average winning probability, one of the player-specific control variables, is not available in 846 choice situations. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Table 19: Regression results for full sample, Logit and Probit specifications

	<i>Dependent variable:</i>					
	Insurance choice dummy					
	Logit			Probit		
	(1)	(2)	(3)	(1)	(2)	(3)
Expected Value	0.011*** (4.894)	0.014*** (5.652)	0.020*** (6.480)	0.008*** (4.032)	0.010*** (4.790)	0.012*** (5.645)
Variance	-0.004* (-1.966)	-0.006* (-2.338)	-0.004 (-1.389)	-0.002 (-1.417)	-0.003 (-1.558)	-0.002 (-1.622)
Skewness	-0.022*** (-16.275)	-0.022*** (-18.065)	-0.023*** (-18.391)	-0.023*** (-16.129)	-0.023*** (-18.170)	-0.023*** (-18.512)
Player-specific controls	No	Yes	Yes	No	Yes	Yes
Hand-specific controls	No	No	Yes	No	No	Yes
Observations	4,450,585	4,449,739	4,449,739	4,450,585	4,449,739	4,449,739

Note: The table reports (average) marginal effects from estimating Logit & Probit specifications according to Equation 5. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. The values of Wald test statistics (for testing the null hypothesis that coefficients are zero) are provided in parentheses, using standard errors clustered at the individual level. Number of observations in Columns 2 and 3 differ because the average winning probability, one of the player-specific control variables, is not available in 846 choice situations. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Table 20: Regression results, employing coefficient of variation as a measure of dispersion instead of expected value and variance

	<i>Dependent variable:</i>			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Coefficient of Variation	0.003** (3.184)	0.003** (3.085)	0.003** (3.133)	0.0001 (0.155)
Skewness	-0.026*** (-17.011)	-0.026*** (-17.168)	-0.026*** (-17.503)	-0.024*** (-16.542)
Constant	0.171*** (41.937)	0.171*** (47.525)	0.171*** (47.798)	
Player-specific controls	No	Yes	Yes	No
Hand-specific controls	No	No	Yes	Yes
Fixed effects	No	No	No	Yes
Observations	4,450,585	4,449,739	4,449,739	4,450,585

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5). The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. Number of observations in Columns 2 and 3 differ because the average winning probability, one of the player-specific control variables, is not available in 846 choice situations. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Table 21: Regression results for a subsample with the net pot to be trimmed between 1%- and 99%-percentiles

	<i>Dependent variable:</i>			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Expected Value	0.009*** (4.778)	0.012*** (6.691)	0.012*** (6.713)	0.024*** (14.446)
Variance	-0.005*** (-5.843)	-0.006*** (-7.162)	-0.006*** (-7.168)	-0.009*** (-11.648)
Skewness	-0.022*** (-16.184)	-0.022*** (-17.231)	-0.022*** (-17.596)	-0.020*** (-16.245)
Constant	0.173*** (42.186)	0.173*** (47.56)	0.173*** (47.835)	
Player-specific controls	No	Yes	Yes	No
Hand-specific controls	No	No	Yes	Yes
Player Fixed effects	No	No	No	Yes
Observations	4,361,753	4,360,925	4,360,925	4,361,753

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5). The underlying sample is trimmed between 1%- and 99%-percentiles of the net pot. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. Number of observations in Columns 2 and 3 differ because the average winning probability, one of the player-specific control variables, is not available in 828 choice situations. *: p<0.1; **: p<0.05; ***: p<0.01.

Table 22: Regression results for the subsample of players that face both left- and right-skewed showdown situations

	<i>Dependent variable:</i>			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Expected Value	0.002 (1.948)	0.003*** (3.36)	0.003*** (3.528)	0.007*** (7.147)
Variance	-0.0003 (-0.889)	-0.001 (-1.883)	-0.001 (-1.556)	-0.003*** (-4.748)
Skewness	-0.023*** (-17.172)	-0.023*** (-18.555)	-0.024*** (-18.977)	-0.023*** (-18.601)
Constant	0.171*** (41.666)	0.171*** (47.257)	0.171*** (47.535)	
Player-specific controls	No	Yes	Yes	No
Hand-specific controls	No	No	Yes	Yes
Fixed effects	No	No	No	Yes
Observations	4,415,012	4,414,920	4,414,920	4,415,012

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5). Compared to Table 3, we exclude all observations of players that do not face at least one right- and one left-skewed lottery. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. Number of observations in Columns 2 and 3 differ because the average winning probability, one of the player-specific control variables, is not available in 92 choice situations. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Table 23: Regression results for a subsample, excluding players who never or always choose the insurance option

	<i>Dependent variable:</i>			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Expected Value	0.010*** (4.621)	0.011*** (5.036)	0.012*** (5.053)	0.017*** (5.734)
Variance	-0.004** (-3.186)	-0.004*** (-3.499)	-0.004*** (-3.445)	-0.007*** (-4.072)
Skewness	-0.030*** (-18.264)	-0.029*** (-18.779)	-0.029*** (-18.998)	-0.028*** (-18.616)
Constant	0.218*** (44.926)	0.218*** (52.943)	0.218*** (53.043)	
Player-specific controls	No	Yes	Yes	No
Hand-specific controls	No	No	Yes	Yes
Fixed effects	No	No	No	Yes
Observations	3,476,736	3,476,692	3,476,692	3,476,736

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5). Compared to Table 3, we exclude all observations of players here who never or always choose the insurance option. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. Number of observations in Columns 2 and 3 differ because the average winning probability, one of the player-specific control variables, is not available in 44 choice situations. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Table 24: Regression results for the subsample in which no split pots occur

	<i>Dependent variable:</i>			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Expected Value	0.002** (2.847)	0.003*** (3.961)	0.004*** (4.219)	0.008*** (7.365)
Variance	−0.001 (−1.308)	−0.001* (−2.062)	−0.001 (−1.630)	−0.003*** (−4.127)
Skewness	−0.024*** (−17.199)	−0.024*** (−18.517)	−0.024*** (−18.947)	−0.024*** (−18.625)
Constant	0.177*** (42.277)	0.177*** (47.952)	0.177*** (48.249)	
Player-specific controls	No	Yes	Yes	No
Hand-specific controls	No	No	Yes	Yes
Fixed effects	No	No	No	Yes
Observations	4,154,930	4,154,123	4,154,123	4,154,930

Note: The table reports OLS regression coefficients of our main empirical specification (Equation 5). Compared to Table 3, we exclude all observations that result in a split pot ex-post. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. The main independent variables of interest are the expected value, variance, and skewness of the underlying lottery. Additionally we add fixed effects for different games (on the stake level) and different months to control for the unobserved heterogeneity across games and over time. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. Number of observations in Columns 2 and 3 differ because the average winning probability, one of the player-specific control variables, is not available in 807 choice situations. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

Table 25: Regression results including the dummy for the insurance choice of the other player and the dummy indicating whether the last action of the respective player was a call

	<i>Dependent variable:</i>					
	Insurance choice dummy					
	(1a)	(1b)	(1c)	(2)	(3)	(4)
Expected Value	0.002* (2.019)	0.002* (2.014)	0.002* (2.017)	0.003*** (3.359)	0.003*** (3.564)	0.007*** (7.167)
Variance	-0.0004 (-0.956)	-0.0004 (-0.951)	-0.0004 (-0.955)	-0.001 (-1.916)	-0.001 (-1.557)	-0.003*** (-4.761)
Skewness	-0.023*** (-17.239)	-0.023*** (-17.137)	-0.023*** (-17.219)	-0.023*** (-18.612)	-0.024*** (-19.101)	-0.023*** (-18.808)
Other player's insurance	-0.0002 (-0.827)	-0.0002 (-0.829)		0.0003 (0.944)	0.001** (2.948)	0.001*** (4.46)
Call dummy	0.001 (1.329)		0.001 (1.33)	0.0002 (0.482)	0.001 (1.916)	0.001*** (4.342)
Constant	0.171*** (41.933)	0.171*** (41.933)	0.171*** (41.933)	0.171*** (47.518)	0.171*** (47.792)	
Player-specific controls	No	No	No	Yes	Yes	No
Hand-specific controls	No	No	No	No	Yes	Yes
Player fixed effects	No	No	No	No	No	Yes
Observations	4,450,585	4,450,585	4,450,585	4,449,739	4,449,739	4,450,585

Note: The table reports OLS regression coefficients for our main empirical specification (Equation 5), including additional controls: the dummy indicating the insurance choice of the other player and the dummy indicating whether the last action of the respective player was a call. Column 1a reports results for the basic specification, including both of these additional dummies, Columns 1b and 1c report results for specifications in which each dummy is included separately. The dependent variable is a binary indicator that equals 1 if a player chooses the insurance option and zero otherwise. All specifications include fixed effects for different games (on the stake level) and different months. The independent variables enter the regression as standardized z-scores. Corresponding t-statistics are provided in parentheses, using standard errors clustered at the individual level. The number of observations in Columns 2 and 3 differs because the average winning probability, one of the player-specific control variables, is not available in 846 choice situations. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.