Skewness Preferences: Evidence from Online Poker

Markus Dertwinkel-Kalt[1], Johannes Kasinger[2] and Dmitrij Schneider[3]

[1]University of Münster
[2]Goethe University Frankfurt & SAFE
[3]Heinrich Heine University Düsseldorf & DICE

August 2022

Motivation

- Standard portfolio theory: mean vs. variance is the main trade-off in choice under risk (Markowitz, 1952)
- Behavioral literature: higher-order moments are important—in particular skewness
- *Skewness preference*: preference for positive and aversion toward negative skewness
- A central prediction of most behavioral models, e.g. prospect theory (Kahnemann & Tversky, 1979), disappointment aversion (Gul, 1991) or salience theory (Bordalo et al. 2012)
- Explanation for various and seemingly unrelated puzzles in behavioral economics and finance

Motivation

Studies on "choice under risk" and "skewness preferences" typically rely on

i) **Laboratory experiments**
   + Clean identification in controlled environment with simple and clear probability distribution
   – Highly stylized with small incentives and sample size
   – Subjects with poor understanding of probabilities and choice at hand

ii) **Real-world decision situations**:
   + Relevant choice situations with large incentives and number of observations
   – Complex underlying probability distribution estimated based on past data
   – Strong assumptions about decision maker's knowledge about relevant probability distribution

This paper

Test for skewness preferences in a setup that combines the advantages of lab experiments and real-word decisions:

- Large and unique panel data set of real-world choices (n=4,450,585) with high incentives (Mean of lotteries' expect value: $62)
- Probability distribution is simple, clear and transparently displayed to the subjects
- 83,219 online poker players face repeated choices between a binary risk and a safe option
- Binary risks are uniquely determined by their first three moments: expected value, variance and skewness
- Individuals with comparable good understanding of risks and probabilities

Preview of results

- Online poker players reveal a strong and robust preference for skewness
- Variance seems to play a minor part in their choice
- Skewness preferences are most pronounced among experienced and losing players
- The effect of skewness remains significant for winning players, in contrast to the variance effect

Background – Decision of interest

Our setup focuses on "showdown" situations in online poker and exploits a novel insurance option:

- The outcome of a showdown cannot be influenced by players and solely depends on the draw of remaining cards $\rightarrow$ **all relevant information is public and probabilities can be easily calculated**
- In a showdown each player faces a binary gamble: **A) losing and receiving a payoff of zero** or **B) winning the pot** (accumulated bets during a hand).
- The newly introduced **"All-in Cashout"** provides a player the option to choose a safe payoff that equals her expected earning in a showdown situation minus a margin of 1%.

$\rightarrow$ Players face choices of the type **"Binary risk vs. safe option"**

Background – Decision environment

Theoretical Background - Binary risks

Any binary risk $L = (x_1, \pi; x_2, 1-\pi)$ with outcomes $x_2 > x_1$ is uniquely determined by its first three moments:

- Expected value: $E = \pi x_1 + (1-\pi)x_2$
- Variance: $V = \pi(1-\pi)(x_2 - x_1)^2$
- Skewness $S = \frac{2\pi - 1}{\sqrt{\pi(1-\pi)}}$

As a result, for binary risks it is possible to vary one moment while fixing the others $\rightarrow$ allows for clean identification for higher order risk preferences

For a binary lottery skewness is unambiguously defined by the lottery's third standardized central moment

Example Skewness Example Mao pairs

Theoretical Background - Skewness Preferences

Conventional wisdom: People face a trade-off between higher expected value and lower variance. This can be overturned by skewness preferences:

people like positively skewed, but dislike negatively skewed risks

Definition 1 (Skewness Preferences)

For any $E \in \mathbb{R}$ *and* $V \in \mathbb{R}_+$*, an agent reveals a preference for skewness if she prefers the binary lottery* $L(E, V, S)$ *over the safe option that pays the lottery's expected value* E *if and only if* S *is sufficiently large.*

In our setup, skewness preferences would predict a **negative relation between insurance choice and the lottery's skewness**, keeping all other moments constant

Data set

- Based on (all) 35,529,631 distinct Omaha poker cash game hands played between January 2020 and June 2021 on Pokerstars
- 4,450,585 observations, where every observation refers to a unique decision by a single player in a two-person showdown situation
- Panel of 83,219 distinct players, making repeated choices with varying probabilities and payouts (moments)
- Raw data from a commercial poker data provider that collects and stores hand histories for Omaha Poker cash games played on Pokerstars

Example raw data

Descriptives: Main variables of interest

Distribution of insurance choices:

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

Summary statistics on the probability moments of insurance choice:

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

For each right-skewed lottery ($\pi < 0.5$), there is exactly one complementary left-skewed lottery ($1 - \pi > 0.5$) with identical variance but inverse skewness

Share of insurance choices for different winning probability ranges

Subjects who face a negatively-skewed lottery choose the insurance option in 20% of cases, while subjects who face a positively-skewed risk do so in only 14% of the cases (p-value < 0.0001)

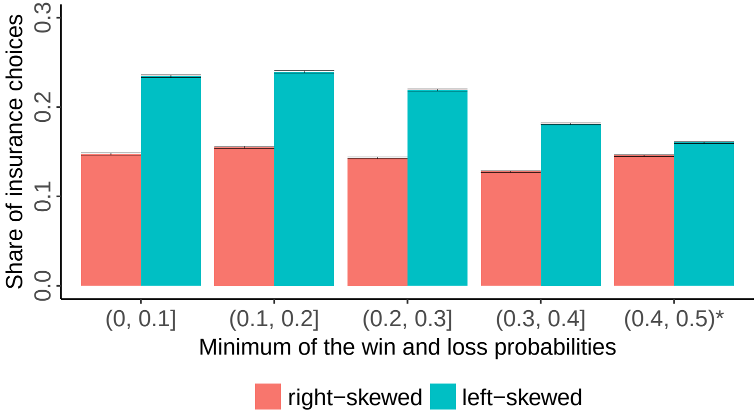

Empirical Strategy

In our main specifications, we follow Mitton & Vorkink (2007) in assuming utility that is linear in the risk moments.
We estimate following reduced-form equation:

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

- y_{ij}: insurance choice dummy
- E_j: expected value
- V_j: variance
- S_j: skewness
- Z_i: player-specific characteristics
- λ_t: time fixed effects
- ψ_z: stake fixed effects
- α_i: player fixed effects

Regression results (LPM) for full sample and standardized variables

	Dependent variable:			
	Insurance choice dummy			
	(1)	(2)	(3)	(4)
Expected value	0.002***	0.003***	0.003***	0.007***
	(8.066)	(12.939)	(13.542)	(7.166)
Variance	−0.0004*	−0.001***	−0.001***	−0.003***
	(−1.689)	(−3.268)	(−2.876)	(−4.757)
Skewness	−0.023***	−0.023***	−0.024***	−0.023***
	(−124.166)	(−124.124)	(−125.872)	(−18.577)
Constant	0.171***	0.171***	0.171***	0.171***
	(965.613)	(974.794)	(966.475)	(3,676.599)
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

Share of insurance choices depending on players' experience

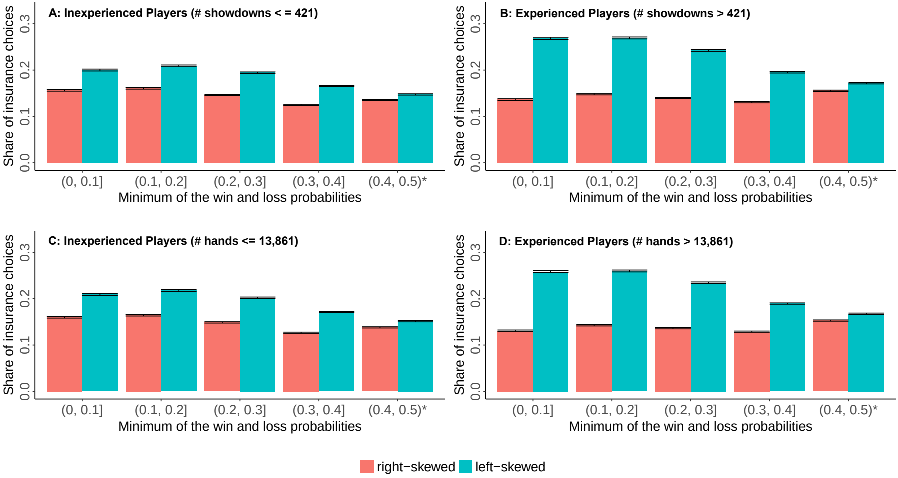

Regression results (LPM) for successful and unsuccessful players

	Dependent variable: 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***	0.002***	0.011***	0.003***	0.002***	0.002***	0.007***	0.002
	(14.164)	(8.058)	(5.599)	(3.343)	(5.591)	(5.586)	(2.690)	(1.636)
Variance	−0.002***	−0.001***	−0.004***	−0.002***	−0.0001	−0.001***	−0.002*	−0.001*
	(−3.605)	(−6.290)	(−3.512)	(−3.088)	(−0.560)	(−5.238)	(1.887)	(−1.690)
Skewness	−0.029***	−0.011***	−0.027***	−0.011***	−0.049***	−0.010***	−0.048***	−0.011***
	(−123.741)	(−40.649)	(−17.422)	(−7.172)	(−111.075)	(−34.107)	(−13.868)	(−5.264)
Constant	0.193***	0.114***	0.193***	0.114***	0.231***	0.105***	0.231***	0.105***
	(879.277)	(410.847)	(2,633.329)	(1,664.233)	(1,462.782)	(340.693)	(2,558.703)	(1,264.016)
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

Robustness Checks

- Excluding all hands that resulted in a split pot (in these cases the underlying risk is not binary)
- Regressions for non-standardized independent variables
- Logit/Probit specifications
- Excluding outliers
- Using the coefficient of variation instead of expected value and variance.
- Considering sample that only includes players that face both types of lotteries – left- and right-skewed – at least once
- Excluding all observations of players who never or always choose the insurance option

Conclusion

- The introduction of the insurance option in online poker allows us to cleanly test for higher order risk preferences in a large and unique set of observational panel data
- We find a strong and robust evidence for skewness preferences, both in qualitative as well as in quantitative terms.
- The effect of skewness is strongest for experienced and losing players but remain significant for winning players
- Variance seems to play a minor part in our sample
- Our results provide important real-world implications for the motivation of individual investors, bettors and entrepreneurs

Additional slides

Share of split pots depending on the expected winning share of the pot

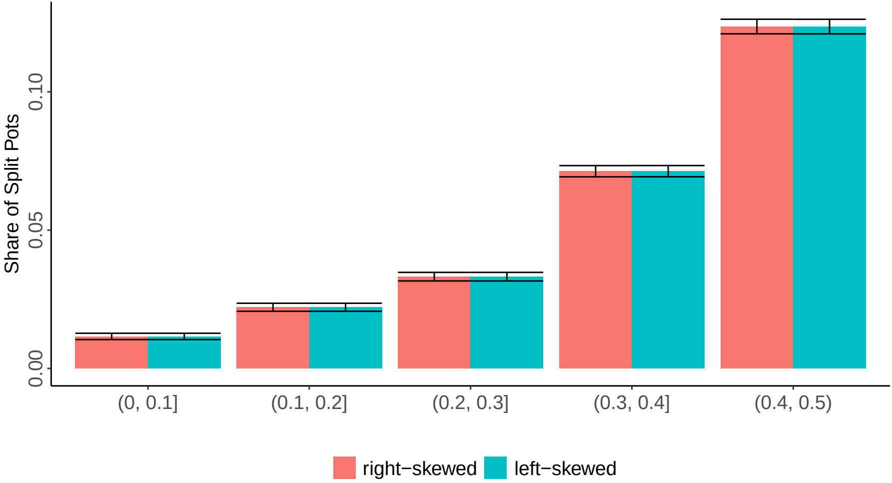

Back

Player- and hand-specific characteristics

- Characteristics of the respective hand, in particular:
  - Stake at the table
  - Stack of each player
  - Position of a player at the table
  - Dummy indicating whether a player risks her entire stack
  - Weekday
- Characteristics of each individual player based on all hands:
  - average winning probability at showdown
  - profit of player per hundred hands
  - number of hands played
  - number of experienced showdown situations

Back

More Descriptives

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Number of hands played	18,846	991.95	3,571.25	1.00	48.00	154.00	565.00	88,896.00
Number of experienced showdown situations	18,846	26.081	75.680	1.00	2.00	6.00	19.00	2,156.00
Average winning probability	18,660	0.45	0.16	0.00	0.37	0.45	0.53	1.00
Profit per hundred hands	18,846	−75.83	1,039.96	−99,020.00	−47.24	−11.76	−0.03	13,060.73

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Stake	491,520	0.631	1.978	0.010	0.100	0.250	0.500	20.000
Stack	491,520	65.022	238.376	0.050	6.400	17.100	46.880	13,863.000

Statistic	N	Risk-all-stack=1	Risk-all-stack=0	Mean	St. Dev.
Risk-All-Stack Dummy	491,520	245,797	245,723	0.500	0.500

	BB	BTN	CO	EP	MP	SB
Frequency	101,961	99,622	85,543	38,884	68,839	96,671

	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Frequency	65,744	68,391	66,766	63,814	74,798	80,843	71,164

Back

Theoretical Background - Skewness Preferences

Most well-known puzzles in choice under risk can be understood as manifestations of **skewness preferences** such as

- the favorite-longshot bias (e.g. Snowberg and Wolfers 2010),
- simultaneous demand for lottery-like gambles and insurance policies (e.g. Kahneman and Tversky 1979, Sydnor 2010, Garett and Sobel 1999),
- Allais paradox (e.g. Allais 1953),

as well as many puzzles in finance and/or labor economics, such as

- the growth puzzle (e.g. Fama and French 1992, Bordalo et al. 2013),
- many instances of portfolio underdiversification (e.g. Mitton and Vorkink 2007),
- abnormal pricing of initial public offerings (e.g. Green and Hwang 2012)

Theoretical Background - Skewness Preferences

For standard utility functions, skewness preferences are **not coherent with EUT**, but can be explained by theories of non-linear probability weighting, such as

- Cumulative Prospect Theory (Kahneman and Tversky 1992),
- Regret & Salience Theory (Loomes and Sugden, 1982; Bordalo et al. 2012),
- Disappointment Aversion (Gul 1991).

The skewness of a probability distribution

Skewness typically refers to the central standardized third moment.

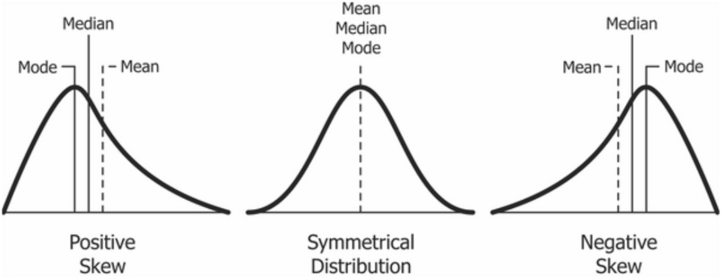

- Right-skewed = positively skewed ($S(L) > 0$): tail on the right side of the probability distribution is long → "large pos. payoff with a small probability."
- Left-skewed = negatively skewed ($S(L) < 0$): tail on the left side of the probability distribution is long → "large neg. payoff with a small probability."

Back

Example of Mao pair

Example of Mao pair – same expected value ($E = 108$) and same variance ($V = 1296$), but different direction of skewness ($S = -S$):

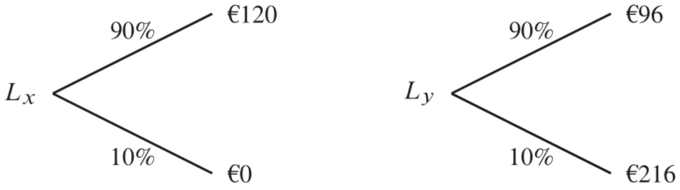

FIGURE 1. An example of a Mao pair.

Source: Dertwinkel-Kalt & Köster (2020)

Back

Data set

Raw data from a commercial poker data provider that collects and stores hand histories for every Omaha Poker cash game played on Pokerstars (and other platforms). Example of a hand history:

```
PokerStars Hand #206477180324:  Hold'em No Limit ($0.01/$0.02 USD) - 2019/11/24 9:21:14 ET
Table 'Adria V' 6-max Seat #2 is the button
Seat 1: peanut_no. 1 ($2.18 in chips)
Seat 2: Narkin55 ($3 in chips)
Seat 3: fablitoA ($0.73 in chips)
Seat 4: netutakogo2 ($1.94 in chips)
Seat 6: Monax0_o ($2.92 in chips)
fablitoA: posts small blind $0.01
netutakogo2: posts big blind $0.02
*** HOLE CARDS ***
Dealt to peanut_no. 1 [Tc 8h]
Monax0_o: folds
peanut_no. 1: raises $2.16 to $2.18 and is all-in
Narkin55: folds
fablitoA: calls $0.72 and is all-in
netutakogo2: calls $1.92 and is all-in
Uncalled bet ($0.24) returned to peanut_no. 1
*** FLOP *** [9h Td Ts]
*** TURN *** [9h Td Ts] [Ac]
*** RIVER *** [9h Td Ts Ac] [Kd]
*** SHOW DOWN ***
netutakogo2: shows [Ks 7s] (two pair, Kings and Tens)
peanut_no. 1: shows [Tc 8h] (three of a kind, Tens)
fablitoA: shows [Jc Ad] (two pair, Aces and Tens)
peanut_no. 1 cashed out the hand for $1.47 | Cash Out Fee $0.01
*** SUMMARY ***
Total pot $4.61 Main pot $2.11. Side pot $2.34. | Rake $0.16
Board [9h Td Ts Ac Kd]
Seat 1: peanut_no. 1 showed [Tc 8h] and won ($4.45) with three of a kind, Tens (pot not awarded as player
cashed out)
Seat 2: Narkin55 (button) folded before Flop (didn't bet)
Seat 3: fablitoA (small blind) showed [Jc Ad] and lost with two pair, Aces and Tens
Seat 4: netutakogo2 (big blind) showed [Ks 7s] and lost with two pair, Kings and Tens
Seat 6: Monax0_o folded before Flop (didn't bet)
```

Limitations

I. Players do not receive the expected value but only 99% of the expected value when they choose the insurance option
II. The risks in the considered showdown situations are not always binary due to the possibility of a split pot.
    - Split pots arise when players hold the same hand after all cards are dealt. In this case, each of the involved players that went to showdown is awarded half of the pot: 6 % of all showdown situations in our sample. Share of Split Pots
    - In case of a split pot possibility we only observe a "weighted winning probability" (expected share of pot)
    - Results are robust if we exclude all hands that result in a split pot: Regression without Split Pots