PNAS RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES

Frequent winners explain apparent skewness preferences in experience-based decisions

Sebastian Olschewski[a,b,1,2], Mikhail S. Spektor[c,d,1,2], and Gaël Le Mens[d,e,f]

Edited by Elke Weber, Princeton University, Princeton, NJ; received October 13, 2023; accepted January 17, 2024

Do people's attitudes toward the (a)symmetry of an outcome distribution affect their choices? Financial investors seek return distributions with frequent small returns but few large ones, consistent with leading models of choice in economics and finance that assume right-skewed preferences. In contrast, many experiments in which decision-makers learn about choice options through experience find the opposite choice tendency, in favor of left-skewed options. To reconcile these seemingly contradicting findings, the present work investigates the effect of skewness on choices in experience-based decisions. Across seven studies, we show that apparent preferences for left-skewed outcome distributions are a consequence of those distributions having a higher value in most direct outcome comparisons, a "frequent-winner effect." By manipulating which option is the frequent winner, we show that choice tendencies for frequent winners can be obtained even with identical outcome distributions. Moreover, systematic choice tendencies in favor of right- or left-skewed options can be obtained by manipulating which option is experienced as the frequent winner. We also find evidence for an intrinsic preference for right-skewed outcome distributions. The frequent-winner phenomenon is robust to variations in outcome distributions and experimental paradigms. These findings are confirmed by computational analyses in which a reinforcement-learning model capturing frequent winning and intrinsic skewness preferences provides the best account of the data. Our work reconciles conflicting findings of aggregated behavior in financial markets and experiments and highlights the need for theories of decision-making sensitive to joint outcome distributions of the available options.

decisions from experience | skewness | financial decision-making | reinforcement learning | higher-order risk preferences

Significance

In contexts such as investment or risky choice, decision-makers seem to prefer options with right-skewed reward distributions that offer high rewards with small probabilities. In contrast, when people experience outcome realizations they show a choice tendency in favor of left-skewed distributions. To reconcile this contradiction, we propose that the cognitive process of comparing individual outcomes with each other is crucial to understand choices. Seemingly contradictory patterns arise because the skewness of a distribution is directly related to the proportion of "wins" in direct outcome comparisons. Our results imply that theories of experience-based decision-making must take joint outcome distributions into account to avoid mistaken conclusions about skewness preferences.

Almost all important decisions in life are made under risk or uncertainty, including career choices, choosing whether or whom to marry, and how much and where to invest pension savings. Understanding behavior in these situations generally entails analyzing how the characteristics of the available options' outcome distributions affect these choices. A vast amount of research has focused on the variability of outcome distributions, but much less research has studied their (a)symmetry, as captured by their skewness. In a nutshell, skewness determines the balance between the frequency and extremeness of outcomes: right-skewed outcome distributions have below-average outcomes most of the time but also some large outcomes, whereas left-skewed outcome distributions have above-average outcomes most of the time but also some small outcomes. Understanding how skewness affects choices is important for researchers and practitioners alike because of the abundance of skewed outcome distributions in high-stakes decisions. Relevant domains include investments in new ventures, where few firms will lead to high returns but most firms will fail (i.e., a right-skewed return distribution), and health-related decisions, where some lifestyle choices (e.g., smoking or drinking alcohol) have relatively rare but very dire health consequences (i.e., a left-skewed outcome distribution).

The current literature does not provide a clear answer regarding the direction of the effect of skewness on choices. On the one hand, some empirical studies and theories argue for a preference for relatively right-skewed outcome distributions (e.g., ref. 1). On the other hand, experiencing the outcome distributions seems to reverse this preference, leading to choice tendencies in favor of relatively left-skewed outcome distributions (e.g., ref. 2). The present research ties together these two seemingly contradictory sets of results by demonstrating how the cognitive process of comparing outcome realizations affects preferential choices.

A large body of literature in economics and finance suggests a preference for right-skewed outcome distributions. For example, lottery-type stocks, referring to stocks with high variance and right skewness in daily returns, have lower expected future returns

Author contributions: S.O., M.S.S., and G.L.M. designed research; S.O. and M.S.S. performed research; S.O. and M.S.S. analyzed data; and S.O., M.S.S., and G.L.M. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]S.O. and M.S.S. contributed equally to this work.

[2]To whom correspondence may be addressed. Email: sebastian.olschewski@unibas.ch or mikhail@spektor.ch.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2317751121/-/DCSupplemental.

Published March 15, 2024.

than alternative stocks, and in particular, individual investors seem to prefer these type of stocks (3, 4). Similarly, investors pay a premium for start-up companies, most of which go bankrupt within the first few years, hoping for the rare occurrence of them becoming "unicorns" (with valuation into the billions)—consistent with a preference for right-skewed distributions of returns (5, 6). Finally, in horse racing, the phenomenon of overvaluing the unlikely winners with high odds (i.e., longshots) has been known for decades (7) and has been linked to skewness preferences (8). Laboratory experiments with humans (9–14) and nonhuman animals (15–17) provide additional evidence for this pattern.

Consistent with these findings, several influential theories of decisions under risk in economics and finance assume implicit or explicit preferences for options with relatively right-skewed outcome distributions (1, 18–21). According to prospect theory (1), the most prominent of these theories, people prefer options with relatively right-skewed outcome distributions because the decision weights they attach to outcomes are non-linear distortions of the objective probabilities. For example, people buy lottery tickets (i.e., right-skewed options) and at the same time insurances (i.e., they dislike the left-skewed option of not being insured) because the unlikely but extreme outcome—the win in the case of a lottery, or the loss in the case of an uninsured catastrophe—receives a higher decision weight than its probability would imply.

In contrast to these observations in economics and finance, experimental research on experience-based decisions has generally found a choice tendency for relatively left-skewed options [(2, 22–25), but see refs. 26 and 27 for limitations]. Similarly, a recent study examining a closely related attitude, prudence, found that experience diminishes the effect of skewness on choices based on experience (28). In this line of research, people observe realizations from the options' outcome distributions and have to base their choices on the inferences they draw from these observations without having any information about the true underlying distributions. Mirroring the dominant explanation for choices of relatively right-skewed outcome distributions, the observed preferences are typically explained by a disproportionally low decision weight of low-probability outcomes (see ref. 29 for an overview).

The discrepancy between the conclusions drawn from research on experience-based decisions and those in financial markets is surprising because decisions in the latter are also based on individual experiences (30, 31). For example, personal experiences with boom-and-bust phases of markets affect stock market participation (32–34) and experienced returns of past investments predict future investments, in line with fundamental principles of reinforcement learning (35). Moreover, the impact of experiences and learning in financial decision-making has been corroborated in laboratory experiments (36–39).

How could this discrepancy be explained? We argue that an important difference between experimental experience-based decisions and investment decisions in naturally occurring environments pertains to how people process outcome information. While experimental designs often emphasize the comparison of realizations from the available options' outcome distributions (e.g., refs. 40–43), information provided to investors about financial assets tends to emphasize individual properties of the assets' past outcome distributions (such as riskiness) or cumulative performance relative to a benchmark (such as a market index). It is possible to compute the frequency of one asset having higher per-period returns than another asset from financial time series, but this information is generally not displayed by default on the platforms people use to make their investment decisions. In other words, this choice environment does not emphasize the comparison of the realizations of the available options' outcome distributions.

Consider a standard experiment in which the participants' task is to choose the preferred option out of two available ones. If realizations from two options are presented next to each other, this encourages decision-makers to compare these realizations with one another. Choosing the option that has the higher outcome (i.e., "wins") in most of these comparisons (henceforth frequent winner) is a straight-forward strategy that minimizes the probability of (anticipated) regret (44, 45). In line with this prediction, the frequency with which one option generated the higher outcome has been found to affect choices in binary experience-based decisions in a variety of settings (25, 46–48).

We use this observation about the frequent-winner effect to explain why, in experience-based decisions between a right-skewed and a left-skewed option, people tend to choose the left-skewed option. If the two distributions have the same mean and variance, the left-skewed option will win most outcome comparisons (see Fig. 1 *B*–*D* for an illustration). For the case of discrete two-outcome distributions, this relationship has been proven to hold (49). In other words: Even if people (weakly) prefer right-skewed outcome distributions, a choice situation that facilitates the pairwise comparison process between individual outcomes can shift their choices in favor of relatively left-skewed options, resulting in an apparent preference for left-skewed options. Therefore, observing that people choose a relatively left-skewed option does not necessarily imply that their choices are driven by an intrinsic preference for left-skewed outcome distribution.

The goal of the present work is to investigate skewness preferences in experience-based decisions, with a focus on the explanatory power of comparison processes between individual realizations of two options. We rely on an experimental paradigm that allows for the precise control of the outcome-comparison process and show how choice tendencies for left-skewed outcome distributions can reverse when the direct comparison of the options' outcomes favors another option. Computational modeling of data from the reported experiments confirms that two factors are necessary to explain the data: an intrinsic preference for right-skewed outcome distributions and a process of direct comparison of options' outcomes.

Results

Study participants played multiple rounds of the "broker game", a variation of the classical decisions-from-experience paradigm (2). Each round contained two "stocks," represented by unlabeled boxes that each produced a sequence of dividends. Participants were not provided with any information about the process that generated the dividends. They obtained information about the stocks by observing 30 dividends from each of them. The dividends were presented simultaneously next to each other in rapid succession (approximately 1 sample every 1.5 s), thereby carefully controlling for the comparison process between the individual realizations of the available options. After observing all outcomes, participants were asked to choose the stock from which they would like to obtain a dividend (Fig. 1*A*). The results of all three main studies are shown in Fig. 2.

Study 1: Skewness and Frequent Winners Are Confounded. In study 1, 99 participants (53 female, 43 male, 3 not disclosed; age: 19 to 50, $M = 25.9$, $SD = 5.9$) encountered each of

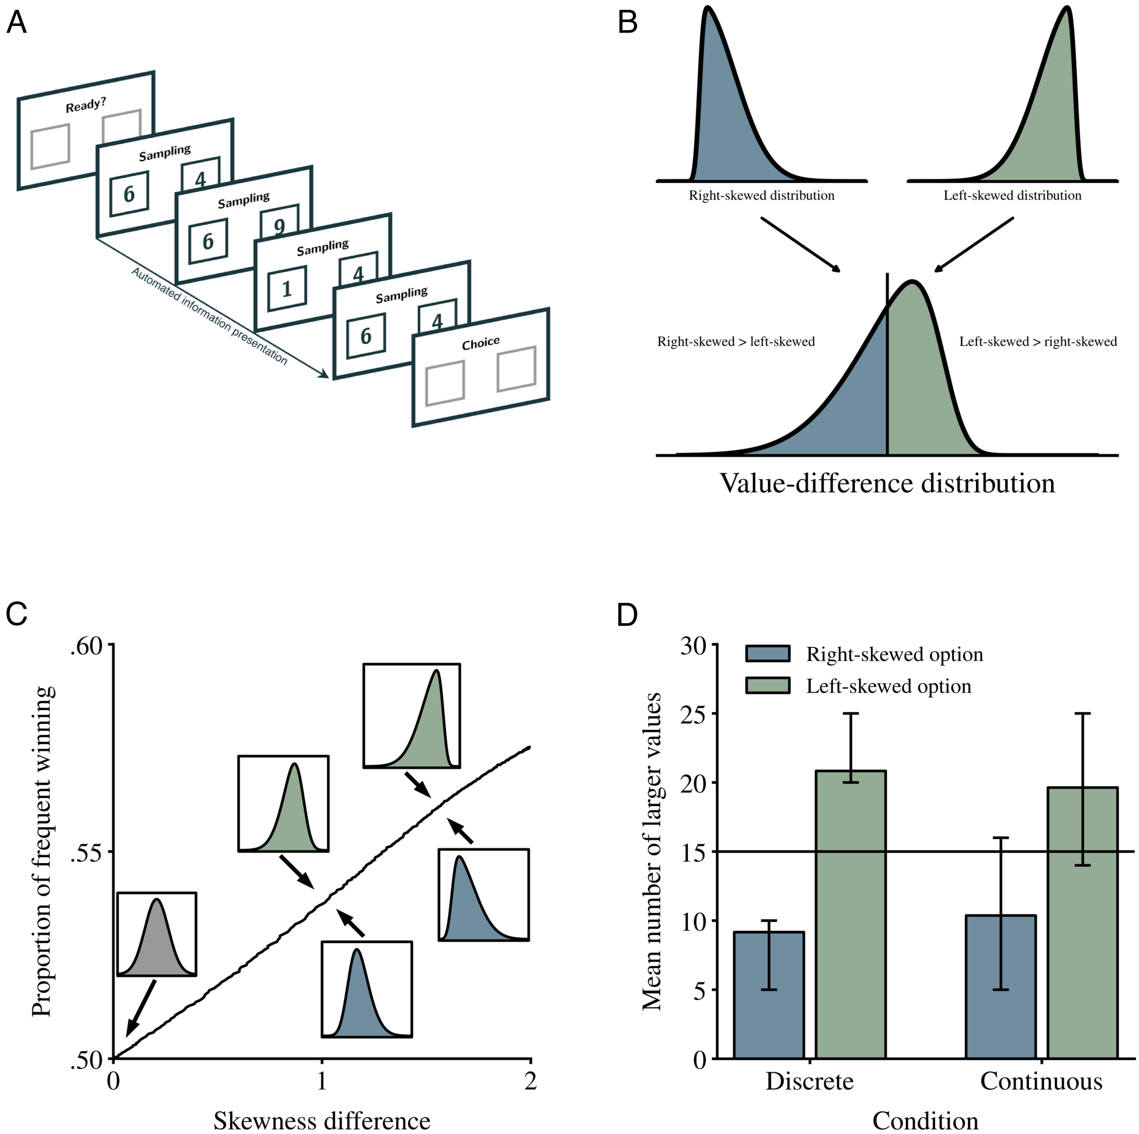

Fig. 1. Main design features. (*A*) illustrates the experimental paradigm in which participants observed a series of outcomes from two outcome distributions and made a single consequential choice afterward. (*B*) illustrates the statistical regularity that, in direct comparison with right-skewed distributions with the same mean and variance, left-skewed distributions have higher values most of the time. The *Bottom* panel shows the expected proportion of times each of the distributions has the higher value in random independent samples. (*C*) illustrates how the statistical frequent-winner effect in (*B*) changes as a function of the relative skewness difference between two options, using skew-normal distributions with standardized shape parameters between 0 and 50 and its inverted mean-matched counterpart. (*D*) shows how the options used in the present experimental design (illustration using study 1) vary in the degree of frequent winning (independent draws of 30 samples each). Error bars illustrate the range across 100,000 simulations.

three experimental conditions that differed in the distributions of the stocks' dividends. In the first condition, participants made four independent decisions between two options with two unique dividends each. In each case, both of the options' outcome distributions had the same expected value and variance, but one of them was right-skewed, whereas the other one was left-skewed. In line with past research using similar (but often not variance-matched) outcome distributions in decision-from-experience (2, 29, 50), people predominantly chose the left-skewed option (M = 60.1%, SD = 29.4%), $t(98) = 3.415$, P = 0.001, Cohen's d = 0.343 [95% CI: (0.140, 0.545)]. To examine whether this choice tendency also occurs with continuous outcome distributions, as they are often encountered in financial decisions, the second condition explored decisions between options with continuous outcome distributions—a setting that is only rarely used in experiments (for exceptions, see refs. 14, 51, and 52). In line with the first condition, participants mostly chose the left-skewed option when outcome distributions were continuous (M = 57.6%, SD = 29.8%), $t(98) = 2.533$, P = 0.013, Cohen's d = 0.255 [95% CI: (0.054, 0.454)].

The third condition was critical to assess the frequent-winner effect. Both options contained identical outcomes derived from a non-skewed symmetrical Gaussian outcome distribution. Here, the presentation order of the pairs of outcomes was manipulated such that one of the two options was the frequent winner: in direct comparison, 20 out of the 30 dividends of one of the options were higher than the dividends of the other option. Even though the sampled marginal outcome distributions of both options were identical, participants had a clear tendency to choose the frequent winner (M = 67.2%, SD = 31.5%), $t(98) = 5.429$, $P < 0.001$, Cohen's d = 0.546 [95% CI: (0.333, 0.756)].

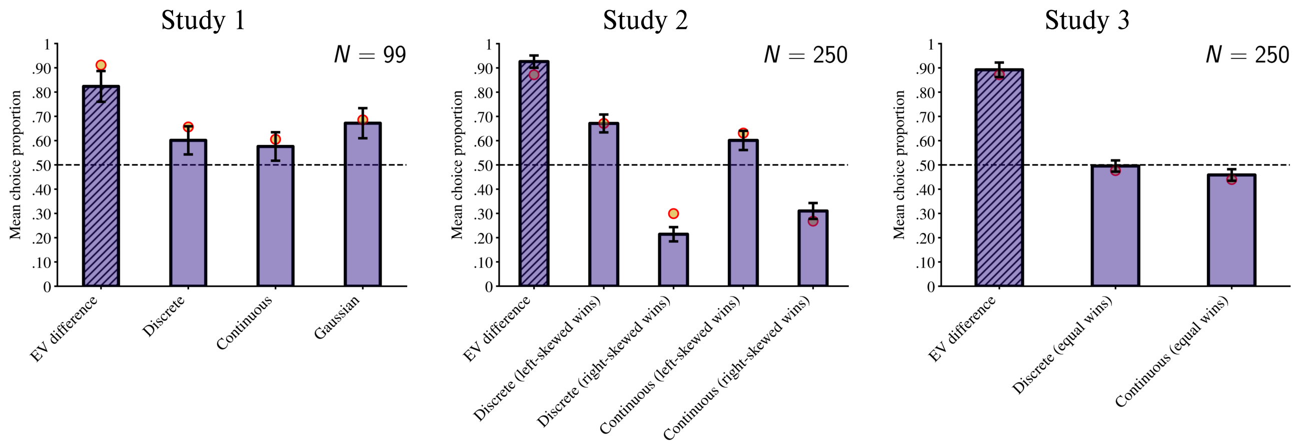

Fig. 2. Behavioral results and model predictions of the three main studies. Each bar corresponds to the mean choice proportion of the left-skewed option. Hatched bars represent choices in which both options were symmetrical and the expected value was different between the options. In this case, the bar height represents choices of the option with the higher expected value. For the "Gaussian" condition, bars represent the choice of the frequently winning option. Dots show the predictions for the winning model that explicitly includes mechanisms to account for the outcome-comparison process and intrinsic skewness preferences ("Tallying RL + Skewness"). The tick labels provide a brief description of the choice situation, see *Materials and Methods* for details. Error bars show the 95% CI of the mean.

The observed choice tendency for the frequent winner in the third condition demonstrates the strong influence of the outcome-comparison process as a determinant of choices. Note that the other two experimental conditions have a naturally occurring frequent winner without additional experimental interventions (Fig. 1*B*). Hence, the process of outcome comparisons could explain the observed choice tendency in favor of relatively left-skewed outcome distributions in experience-based risky choices.

Study 2: Disentangling Skewness from Frequent Winners. Study 2 aims to provide direct evidence for the hypothesis that the process of outcome comparisons explains the effect of skewness on choice by manipulating the option that is the frequent winner without manipulating skewness. Two hundred and fifty participants (149 female, 100 male, 1 not disclosed; age: 19 to 74, $M = 39.2$, $SD = 12.3$) encountered two pairs of experimental conditions involving four choices each. In the first pair of conditions, they made a choice between two options with outcome distributions that had three unique dividends with equal probabilities of occurrence. They had the same expected values and variances and differed only in skewness, with one being left-skewed and the other being right-skewed. In one of the two conditions, the left-skewed option was the frequent winner, and participants were more likely to choose this option than the frequent loser ($M = 67.1\%$, $SD = 29.7\%$), $t(249) = 9.096$, $P < 0.001$, Cohen's $d = 0.575$ (95% CI: [0.441, 0.709]). In the other condition, the right-skewed option was the frequent winner, reversing participants' choice tendency (left-skewed $M = 21.4\%$, $SD = 23.8\%$), $t(249) = -19.036$, $P < 0.001$, Cohen's $d = -1.204$ [95% CI: (−1.366, −1.040)].

The second pair of conditions aimed at examining the robustness of the observed pattern with continuous outcome distributions. Consistent with the results obtained with discrete outcome distributions, participants predominantly chose the left-skewed option when it was the frequent winner ($M = 60.1\%$, $SD = 31.9\%$), $t(249) = 5.001$, $P < 0.001$, Cohen's $d = 0.316$ [95% CI: (0.189, 0.443)], but the choice tendency reversed when the right-skewed option was the frequent winner (left-skewed $M = 31.0\%$, $SD = 26.5\%$), $t(249) = -11.340$, $P < 0.001$, Cohen's $d = -0.717$ [95% CI: (−0.856, −0.578)].

In addition to the hypothesized main effects, a 2 (outcome distribution: continuous vs. discrete) × 2 (frequent winner) repeated-measures ANOVA revealed two noteworthy observations: First, an interaction effect was present such that the effect of the frequent winner was stronger for discrete outcome distributions than for the continuous ones, $\eta_p^2 = 0.144$. Second, controlling for the main effects and their interaction, a substantial deviation from indifference between the right- and left-skewed options arose, suggesting an intrinsic preference for right-skewed options, $\eta_p^2 = 0.123$.

Study 3: Skewness in the Absence of Frequent Winners. In study 3, we aimed to directly test for skewness preferences in the absence of a frequent winner. Two hundred and fifty participants (125 female, 121 male, 4 not disclosed; age: 18 to 79, $M = 39.0$, $SD = 12.9$) made eight decisions in each of two experimental conditions, one containing discrete outcomes and the other continuous outcomes. Neither of the two options was a frequent winner (i.e., each option had higher dividends 50% of the time). The marginal distributions were identical to those of study 2, one right-skewed and one left-skewed distribution. With discrete outcomes, participants were indifferent between the two options (left-skewed $M = 49.6\%$, $SD = 18.6\%$), $t(249) = -0.383$, $P = 0.702$, Cohen's $d = -0.024$ [95% CI: (−0.148, 0.100)]. With continuous outcomes, however, participants tended to choose the option with the right-skewed outcome distribution (left-skewed $M = 45.8\%$, $SD = 19.2\%$), $t(249) = -3.423$, $P = 0.001$, Cohen's $d = -0.217$ [95% CI: (−0.342, −0.091)]. The results of this study provide partial support for the hypothesis that participants have an intrinsic preference for right-skewed outcome distributions.

A Model of Skewness Preferences and the Outcome-Comparison Process. To uncover the mechanisms that best describe behavior in our data, we formalized the two phenomena identified in the above sections (i.e., direct outcome comparison and an intrinsic preference for right-skewed outcome distributions) within a reinforcement-learning framework. Reinforcement learning (54) is arguably the most popular framework to analyze human decision-making based on experience. Within this framework, decision-makers keep track of all options' expected rewards and update their expectations as they observe new realizations from the options' outcome distributions. If the new realizations are higher than what the decision-maker expects, then they increase their reward expectations and vice versa. A core feature of this framework is that the reward expectations are processed for each option independently from the other available options (but see ref. 41 for recent theoretical developments).

We capture the frequent-winner effect by implementing the direct outcome comparison process in the reinforcement learning framework. We implement a tallying mechanism which compares the outcomes of two options by ignoring the magnitude of the difference, such that the subjective reward merely reflects which outcome was "better" (see refs. 55 and 56). We capture a possible intrinsic preference for skewed outcome distributions by boosting the subjective valuation of an option with an additive term proportional to the skewness of the observed outcomes, similar to mean–variance–skewness models of financial decision-making (57, 58).

We compared a random-guessing model ("Guessing") and the standard reinforcement-learning model ("Standard RL") with reinforcement-learning models incorporating the tallying mechanism only ("Tallying RL"), the intrinsic preference of skewness only ("Standard RL + Skewness") and a combination of tallying and intrinsic preference for skewness ("Tallying RL + Skewness"). We estimated the model parameters on the behavioral data from the studies via maximum likelihood. Details are provided in *Materials and Methods* and the resulting parameter estimates are reported in Table 1.

The standard reinforcement-learning model provides a strong baseline, as it is in principle able to capture frequent-winner effects. It can do so by forming reward expectations based on the most recently observed outcomes (captured by setting the learning rate parameter close to its maximum value of 1) (25). In contrast, the best fitting parameters of the standard reinforcement-learning model imply only a low level of recency (learning rate $\alpha = 0.04$, Table 1), consistent with past studies (29, 41, 59). This low level of recency is not able to capture the observed choice proportions in favor of frequently winning options, in particular when both options have identical Gaussian outcome distributions (in study 1) and when choices of skewed options follow the frequent winner (study 2). In other words, a reinforcement-learning model that integrates information within each option (independently from the other options) cannot capture the main phenomena reported here. Nevertheless, this model ("Standard RL," BIC = 13,539) captured some behavioral patterns and was better than a guessing model that predicts 50–50 choices on all trials ("Guessing," BIC = 14,398, BF = 2.82×10^{186}). Adding a parameter for skewness preferences results in a slightly better complexity-corrected fit ("Standard RL + Skewness," BIC = 13,535) than the standard reinforcement-learning model (BF = 8.63). In contrast, a model that keeps track of which option has the better dividend in direct comparisons (ignoring the magnitude of the difference) improves the model fit to a greater degree (Tallying RL, BIC = 12,957, BF = 3.21×10^{126}). Importantly, adding skewness preferences to this model yet again increases the model's performance, capturing all qualitative choice patterns across the three main studies (Tallying RL + Skewness, BIC = 12,894, BF = 5.49×10^{13}).

Table 1. Estimated parameters and model fits

Model	α	θ	ξ	BIC
Guessing	–	–	–	14,398
Standard RL	0.04	34.27	–	13,539
Standard RL + Skewness	0.04	34.29	0.001	13,535
Tallying RL	0.02	2.26	–	12,957
Tallying RL + Skewness	0.02	2.34	0.03	12,894

Guessing = a baseline model that assumes random guessing on each trial. Standard RL = standard reinforcement-learning model. Tallying RL = the RL model with a tallying-outcome-comparison mechanism. Standard RL + Skewness = the standard RL model with intrinsic skewness preferences. Tallying RL + Skewness = the RL model with a tallying mechanism and intrinsic skewness preferences. α = rate of updating of reward expectations (learning rate). θ = choice-sensitivity parameter of the choice rule. ξ = parameter determining the extent of skewness preferences. BIC = Bayesian information criterion (53). See *Materials and Methods* for details.

Studies 4 to 7: Robustness Checks and Extensions. In *SI Appendix*, Fig. S1, we report four additional studies that extend those reported in the main text and highlight the robustness of the frequent-winner phenomenon. More specifically, we show that this phenomenon occurs even with easy-to-process numbers (study 4) and that it does not override individuals' sensitivity to differences in expected values (studies 4 and 5). It occurs in another commonly used paradigm, the multi-armed bandit task (studies 5 to 7), in choice sets with more than two options (studies 6 and 7), and can even lead to reversals of preferences (study 7).

Discussion

The present study investigated people's skewness preferences in experience-based choices, aiming to distinguish between intrinsic skewness preferences and choice patterns arising from a direct comparison of individual outcomes with each other. In a series of experiments, we found that intrinsic preferences for right-skewed outcome distributions can be overridden if another option is a frequent winner—if it has higher values than the other option most of the time. This phenomenon took place across a variety of choice situations and in two different experience-based tasks.

Our results reconcile conflicting findings from two streams of research. The first stream relies on experimental studies of experience-based decisions and found choice patterns seemingly consistent with a preference for relatively left-skewed outcome distributions (e.g., ref. 2). Our studies show that such choice patterns need not be driven by intrinsic preferences for left-skewed outcome distributions, but instead are entirely compatible with a comparison process in which decision-makers simply track which option provides the highest rewards. Because of the statistical regularity that relatively left-skewed outcome distributions have higher values most of the time (compared to relatively right-skewed outcome distributions), such a comparison process favors left-skewed options. Only in the case in which there is no frequent winner can intrinsic skewness preferences be directly measured. In such setting, participants preferred right-skewed outcome distributions. Therefore, after controlling for the local comparison process of individual realizations, preferences in experience-based experimental settings are consistent with the preferences assumed in the other stream of research that analyzed investment decisions in financial markets (e.g., ref. 35), thereby reconciling these two sets of findings.

Modeling Implications. Models assuming that, in the first stage, information is integrated within options (independently of the other options) and, in the second stage, the integrated values are compared (including standard cases of reinforcement-learning models), are generally unable to predict the frequent-winner phenomena in our data. Rather, because of the predictive accuracy of the Tallying RL + Skewness model, we conclude that individuals take joint outcome distributions into account. Tallying assumes that people engage in ordinal comparisons, ignoring the magnitude of differences. This cognitive process originated in the categorization literature (60) and has successfully been applied in a variety of judgments (e.g., ref. 61), including experience-based choices (55, 56). This theory-driven and broadly applicable cognitive mechanism is able to explain the frequent-winner effects observed in our studies. Alternative theoretical approaches are also able to capture frequent-winner effects. First, it could result from the utility function of a person pursuing the goal of minimizing the probability of (anticipated) regret (44–46); second, it could be a consequence of people making decisions based on a small sample of experienced instances (47, 50, 62–64); third, it could be the result of selective integration of information (42, 43). As can be seen in *SI Appendix*, Fig. S2, a model that relies on selective integration (BIC = 12,962) makes predictions that are virtually indistinguishable from the Tallying RL model, the latter being better than the former by having one fewer free parameter (BF = 15.55).

Importantly, these alternative theoretical approaches account only for the frequent-winner effect in our data. However, in their current forms and to the best of our knowledge, none of these approaches can account for the two additional behavioral phenomena we observed, namely the stronger choice tendency for the frequent winner when the frequent winner is the right-skewed rather than the left-skewed option, and the tendency to choose right-skewed options in the absence of a frequently winning option. The Tallying RL + Skewness model captures these phenomena by assuming intrinsic preferences for right-skewed outcome distributions, similar to mean–variance–skewness models often used in finance (57, 58). Future research could explore augmenting the above-mentioned theoretical approaches with such a component as well. To ultimately determine which approach provides the best quantitative account of behavior would require a more heterogeneous set of stimuli in which the choice options vary not only in their skewness and frequent winning, but also in their variance, expected values, and potentially other features.

The Adaptive Character of the Two Mechanisms. The two mechanisms that were necessary to account for the present data, frequent winning and intrinsic skewness preferences, are adaptive in certain environments. First, relying on frequent winning as a decision-making strategy not only simplifies the decision problem but maximizes expected value when rewards are independently drawn from symmetrical outcome distributions. However, in decision environments where asymmetric or correlated outcome distributions are common, as is often the case in experimental research and in financial markets, choosing the frequent winner can be sub-optimal with respect to an expected-value maximization goal (65). Despite this caveat, choosing the frequent winner can be an evolutionary stable strategy whenever the quality of a decision is evaluated in terms of its outcome, rather than in terms of the decision process (i.e., an outcome bias as described in ref. 66). As an example, if voters assess politicians only based on the outcomes of their policies, these politicians might be tempted to implement policies with a large probability of success (and a small societal impact) rather than a large societal impact (but low probability of success).

Second, preferences for right-skewed outcome distributions can be adaptive as well. This is the case because these preferences can themselves be driven by the disproportionate weighting of extreme values of outcome distributions. Recent work argues that such overweighting of extreme values can be an adaptive information-integration strategy for agents with specific cognitive limitations (e.g., refs. 67–70). These limitations include capacity constraints when sampling events from memory (69) and noise in the information-integration process during decision-making (67). In both of these cases, overweighting extreme events leads to more consistent choices and reduces the risk of neglecting highly impactful events. Yet, even though overweighting extreme events can be adaptive, in our experiments, its impact on choices was overshadowed by the frequent-winner effect. This suggests that overweighting of extreme values will only affect choice patterns when no clear frequent winner can be identified (e.g., because of imprecise perception as in ref. 48).

Limitations and Extensions. It is noteworthy that a clear choice tendency for right-skewed outcome distributions (after controlling for the frequent winner) was only present when the outcome distributions were continuous, but not when they were discrete with three equally probable outcomes. This could result from the fact that the skewness manipulation was less strong in the latter than in the former case (study 3, skewnesses of discrete distributions: ± 0.70, skewnesses of continuous distributions: ± 1.72). This difference was unavoidable because, for three equiprobable outcomes, as used in the discrete condition of study 3, the skewness is bounded in the interval $[-0.71, +0.71]$ (see Proposition 1 in *SI Appendix*). As such, the experimental manipulation was attenuated in the discrete distribution setting, limiting the power to detect intrinsic skewness preferences in this setting.

A possible extension of the present study would involve decisions with more than three options. Such a design would be able to clarify whether frequent-winner effects decrease or increase in more complex choice situations (e.g., choices involving four options, ref. 71). After all, decision-makers cannot reasonably pay attention to all outcomes and need to allocate their attentional resources somehow (72). The computational benefits of simple heuristics (such as paying attention to which option yields the highest outcome on a given trial) become more apparent, but at the same time, the costs of making a sub-optimal choice can increase as the number of potentially better options increases. Past research has already illustrated that people are able to simplify the decision problem by extracting relevant environmental characteristics (73); paying attention to frequent winners might be one way to achieve such a simplification.

Conclusion. Our study elucidates the intricate relationship between skewness preferences and the frequent-winner effect in experience-based choices, underscoring the influence of task environment characteristics on choice tendencies. Our findings reveal that, when the task environment promotes outcome comparisons, an inherent preference for right-skewed options may not necessarily translate into a propensity to select such an option. In fact, choices tend to favor options with left-skewed payoff distributions in these circumstances. As we move forward, a deeper understanding of how skewness impacts choices in specific settings warrants a thorough examination of how the structure of the choice environment affects information processing pertaining to the available options. This approach can contribute to the refinement of decision-making models and the development of more effective interventions and policies to align choice proclivities with underlying preferences.

Materials and Methods

Participants and Procedure. All data of the studies reported in the main text were collected through Prolific, an online crowd-sourcing platform. All participants' first language was English, their approval rate was above 80%, and each individual participated in exactly one of the studies reported here. Sample sizes were determined prior to data collection and all analyses were conducted with all observations without any exclusions. Study 1 aimed to collect data from 100 participants but, due to a technical error, ended up with 99 participants. The other studies aimed to collect data from 250 participants each. The demographics of each study are reported in the main text. Each study took about 15 min to complete and included a flat participation fee of £2.00. Additionally, participants earned a bonus depending on one of their choices in the experiment. In study 1, participants earned an average bonus of £1.11 (range: £0.78 to £1.40). In study 2, the mean was £1.06 (range: £0.80 to £1.32), and in study 3 it was £0.99 (range: £0.48 to £1.49). Ethical approval was obtained from the Warwick Humanities and Social Sciences Research Ethics Committee (HSSREC 225/20–21) under the DR@W agreement (submission ID 542119249).

After giving informed consent, individuals read the instructions at their own pace. The studies were framed as a "broker game" and participants were informed that they would observe 30 past dividends from each of two stocks presented simultaneously to learn about what dividends they offer. It was further explained that after observing the samples, participants would decide from which stock they want to draw a dividend. Sampling started after participants pressed a button and the 30 dividends for each option were presented automatically within two unlabeled rectangles. The dividends of both options were presented in quick succession with an on-screen time of 1,250 ms per sample and a 200 ms inter-sample interval. After all the samples had been displayed, participants were prompted to choose one of the two options either by clicking on the respective rectangle or by pressing a button on the keyboard. All outcome feedback was relegated to the end of the experiment, where one trial was randomly selected and the respective choice in that trial was played out for a monetary bonus.

Design and Stimuli. Each study involved multiple decisions from the participants that were grouped into experimental conditions based on theoretically meaningful properties. In all cases, individuals observed 30 draws from each of the options' outcome distributions.

Study 1. Study 1 contained a total of 14 decisions per participant. Two of these decisions involved two options, one of which was clearly superior to the other. These decisions were used to check for people's sensitivity to expected values, and a vast majority of participants chose the better option (82%). The remaining 12 were divided into three conditions of four decisions each. The four decisions were shifted versions of one another in which both options had identical means (375, 425, 475, or 525 points) and SD ($SD = 20$). Therefore, we describe the main properties of one of the four decisions within each experimental condition, and the other three decisions were shifted versions of it.

The first experimental condition involved two options with discrete outcome distributions and two unique outcomes each, occurring with probabilities of 5/30 and 25/30, respectively. The difference between the two options was that one was right-skewed (standardized skewness of $+1.79$) and the other was left-skewed (standardized skewness of -1.79). We obtained the right-skewed option by randomly drawing an outcome from the mean ± 49 and matching the other outcome to obtain the mean. The left-skewed option was then obtained by mirroring the right-skewed option and matching its mean.

The second experimental condition contained continuous outcome distributions. We created the right-skewed option by drawing random samples from a Gamma distribution with a shape parameter of 0.5 and a scale parameter of 28.28, with the additional condition that there were exactly two extreme events (which we defined as numbers at least 3σ away from the theoretical mean of the distribution). The resulting samples (standardized skewness of $+2.06$) were then linearly transformed to conform with the target mean and SD. The sequences for the left-skewed option were obtained using a similar procedure using draws from a mirrored gamma distribution.

In the third experimental condition, both options contained identical outcomes drawn from a symmetrical Gaussian distribution. The outcomes were random samples from the theoretical distribution, controlling for representativeness of the underlying distribution.

The order with which the outcomes were presented (sequences) was pseudo-randomized to fulfill the following properties: In the experimental conditions 1 and 2, the rare events occurred within bins so that they were roughly equally distributed within the sequences (e.g., in the case of two rare events, one event happened in each half of the sequence). In the experimental condition 3, outcomes of the two identical Gaussian distributions were manipulated in such a way that one of them had the higher outcome in 20 out of 30 samples that were presented together.

Study 2. Study 2 contained a total of 18 decisions. As in study 1, two of those decisions were expected-value-sensitivity trials with a very high accuracy (93%). The remaining 16 trials contained two options with identical means (350, 400, 450, or 500 points) and SD ($SD = 33.69$ for trials with continuous outcome distributions and $SD = 15.15$ for discrete outcome distributions). Eight out of these trials contained options with continuous outcome distributions and a skewness of either $+1.72$ or -1.72. The other eight trials contained options with discrete outcome distributions and a skewness of either $+0.56$ or -0.56.

In contrast to study 1, the outcomes were generated by taking the values of 30 equally spaced quantiles of the theoretical outcome distributions. Additionally, the discrete outcome distributions contained three unique outcomes, each occurring 10 times. The outcomes were manipulated in such a way that in four of the eight trials in each pair of conditions (continuous and discrete outcome distributions), the left-skewed outcome distributions had higher values in 20 out of 30 observations. In the other four trials, the right-skewed outcome distribution had higher values in 20 out of 30 observations.

Study 3. In study 3, there were a total of 18 trials, of which two were used as expected-value-sensitivity checks (89% accuracy). The remaining 16 trials were, much like in study 2, split into eight trials with continuous outcome distributions and eight trials with discrete outcome distributions. The general properties of the outcome distributions were comparable to study 2, apart from the following changes: The SD were identical for continuous and discrete outcome distributions ($SD \approx 33.50$), and the skewness of the discrete outcome distributions was slightly larger than in study 2 (either $+0.70$ or -0.70). The crucial difference to the other studies presented here was that we manipulated the outcomes such that each of the two options had higher values on exactly half of the observations.

Computational Modeling. To provide a model-based characterization of behavior in our studies, we relied on the reinforcement-learning framework (54), one of the standard ways to analyze human learning and decision-making. Within this framework, we used the Q-learning/δ-rule combination as popularized by Rescorla and Wagner (74). The fundamental idea is that decision-makers keep track of a subjective reward expectation for option i on trial t, $Q_{i,t}$, and update this expectation as a fraction α of the reward-prediction error δ (i.e., the difference between the obtained reward, $R_{i,t}$, and the expected one). Formally, the subjective expectation for option i on the next trial $t+1$ is given by

$$Q_{i,t+1} = Q_{i,t} + \alpha \delta_{i,t},$$

where $R_{i,t}$ is the obtained reward, $\delta_{i,t} = (R_{i,t} - Q_{i,t})$, and the initial expectation $Q_{i,0} = 0$. Different values of the learning rate α determine how much decision-makers pay attention to more recent observations. In the most extreme cases, reward expectations are based entirely on the most recent observation when $\alpha = 1$ and are based on the initial expectation when $\alpha = 0$. The reward expectations are then transformed into choice probabilities $P_{i,t}$ using a softmax/multinomial logit choice function:

$$P_{i,t} = \frac{e^{\theta Q_{i,t}}}{\sum_{j=1}^{n} e^{\theta Q_{j,t}}},$$

where n is the number of available options and θ is the choice-sensitivity parameter, such that, with $\theta = 0$, choices are completely random, and become deterministic as θ approaches infinity. Since the present studies do not rely on repeated choices within one trial (except for studies 5 to 7), choice probabilities for each trial were derived based on all 30 dividends from the stocks (i.e., $t = 30$).

To test which mechanisms provide the best account of the choice data, we endowed the standard model with additional mechanisms. The first mechanism, *tallying*, recodes the observed reward into a dummy variable, depending on whether the reward is highest on a given trial or not:

$$R_{i,t} = \begin{cases} 1 & \text{if } R_{i,t} = \max R_{\bullet,t} \\ 0 & \text{otherwise.} \end{cases}$$

Tallying is able to account for frequent-winning preferences by ignoring all differences in reward magnitudes without introducing any new free parameters. To capture the second mechanism of intrinsic skewness preferences, the choice function was adjusted. In particular, we assume that choices are not only driven by $Q_{i,31}$ but also by the observed skewness of options' outcome distributions Skew(i, 31):

$$P_i = \frac{e^{\theta(Q_{i,31}+\xi\times \text{Skew}(i,31))}}{\sum_{j=1}^{n} e^{\theta(Q_{j,31}+\xi\times \text{Skew}(j,31))}}.$$

The parameter ξ determines whether decision-makers prefer right-skewed outcome distributions (when $\xi > 0$), left-skewed outcome distributions (when $\xi < 0$), or are indifferent with respect to skewness (when $\xi = 0$).

In addition to reinforcement learning, we also examined selective integration as an alternative model class. It is popular in decisions from experience and has been successfully applied to explain the effects of variance, task framing, and context on decisions (42, 43). Using the same notation as above, the selective-integration model uses the following accumulation function to calculate the subjective reward expectations:

$$Q_{i,t+1} = (1-\lambda)\times Q_{i,t} + \eta(R_{i,t}, R_{j,t})\times R_{i,t},$$

where λ captures the decaying influence of past observations and η is the selective-integration function that depends on the observed rewards of both options i and j. We use the functional form for η described as model "MD2" in Tsetsos (75):

$$\eta(R_{i,t}, R_{j,t}) = \frac{1}{1+e^{-\beta\times(R_{i,t}-R_{j,t})}},$$

where parameter β governs the effect of the local comparison on the weighting of the observed reward for a given option. When $\beta = 0$, the local comparison does not distort the weighting of the observed rewards between competing options and the model reduces to a context-free learning model with memory decay. As β increases, rewards of the locally winning options get a relatively higher weight in the accumulation process, and other options' rewards are proportionally discounted. The same formulæ with interchanged indices i and j are applied to calculate the reward expectations for option j. The model relies on the same softmax function as the reinforcement-learning models to derive choice probabilities from the subjective reward expectations.

In total six models entered our comparative analysis: a random-guessing model (Guessing), the standard reinforcement-learning model (Standard RL) and reinforcement-learning models incorporating the tallying mechanism only (Tallying RL), the intrinsic preference of skewness only (Standard RL + Skewness), a combination of tallying and intrinsic preference for skewness (Tallying RL + Skewness), and the selective-integration model described above.

To evaluate model performance, we used a complete-pooling maximum-likelihood approach across studies 1 to 3. Every choice from every study was assumed to stem from a model with the same parameter values across participants. While this approach ignores individual differences, it forces models to explain behavioral patterns in all experimental conditions, not only those within a single study. The maximum-likelihood estimates were obtained using a differential-evolution algorithm with a population size of 100. The models were compared using the Bayesian information criterion (BIC; ref. 53) which introduces the idea of parsimony into model comparisons by penalizing models with more free parameters. The degree to which a model is better than another, we relied on the BIC-based Bayes-factor (BF) approximation (76). To ensure comparability of the parameter estimates across models including tallying and those without, all Q values were normalized to sum up to 1 within a trial. Results without this normalization are qualitatively and quantitatively virtually identical and are not reported here.

Data, Materials, and Software Availability. Anonymized data and the full analysis code can be found on the Open Science Framework: https://doi.org/10.17605/OSF.IO/AQJDZ (77).

ACKNOWLEDGMENTS. M.S. and G.L.M. gratefully acknowledge the financial support from the Banco Bilbao Vizcaya Argentaria Foundation (project G999088Q), grant PID2019-105249GBI00/AEI/10.13039/501100011033 from the Spanish Ministerio de Ciencia, Innovacion y Universidades and the Agencia Estatal de Investigacion, and the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2019-000915-S), funded by MCIN/AEI/10.13039/501100011033 and European Regional Development Fund. M.S. was also supported by the Spanish Ministry of Economic Affairs and Digital Transformation through the Juan de la Cierva fellowship (FJC2019–040970–I). G.L.M. also received financial support from European Research Council (ERC) Consolidator Grant 772268 from the European Commission and from a Catalan Institution for Research and Advanced Studies (ICREA) Academia grant. We thank Doron Cohen, Ido Erev, Robin Hogarth, Tomás Lejarraga, Stefano Palminteri, and Konstantinos Tsetsos for their valuable comments and discussions.

Author affiliations: [a]Department of Psychology, University of Basel, 4055 Basel, Switzerland; [b]Warwick Business School, University of Warwick, CV4 7EQ Coventry, United Kingdom; [c]Department of Psychology, University of Warwick, CV4 7EQ Coventry, United Kingdom; [d]Department of Economics and Business, Universitat Pompeu Fabra, 08005 Barcelona, Spain; [e]Barcelona School of Economics (BSE), Barcelona 08005, Spain; and [f]Universitat Pompeu Fabra–Barcelona School of Management, 08008 Barcelona, Spain

1. D. Kahneman, A. Tversky, Prospect theory: An analysis of decision under risk. *Econometrica* **47**, 263–292 (1979).
2. R. Hertwig, G. Barron, E. U. Weber, I. Erev, Decisions from experience and the effect of rare events in risky choice. *Psychol. Sci.* **15**, 534–539 (2004).
3. T. G. Bali, N. Cakici, R. F. Whitelaw, Maxing out: Stocks as lotteries and the cross-section of expected returns. *J. Financ. Econ.* **99**, 427–446 (2011).
4. A. Kumar, Who gambles in the stock market? *J. Fin.* **64**, 1889–1933 (2009).
5. B. Eraker, M. Ready, Do investors overpay for stocks with lottery-like payoffs? An examination of the returns of OTC stocks. *J. Fin. Econ.* **115**, 486–504 (2015).
6. T. C. Green, B. H. Hwang, Initial public offerings as lotteries: Skewness preference and first-day returns. *Manag. Sci.* **58**, 432–444 (2012).
7. R. M. Griffith, Odds adjustments by American horse-race bettors. *Am. J. Psychol.* **62**, 290 (1949).
8. J. Golec, M. Tamarkin, Bettors love skewness, not risk, at the horse track. *J. Polit. Econ.* **106**, 205–225 (1998).
9. P. J. Grossman, C. C. Eckel, Loving the long shot: Risk taking with skewed lotteries. *J. Risk Uncertain.* **51**, 195–217 (2015).
10. C. C. Wu, P. Bossaerts, B. Knutson, The affective impact of financial skewness on neural activity and choice. *PLoS ONE* **6**, 1–7 (2011).
11. S. Ebert, D. Wiesen, Testing for prudence and skewness seeking. *Manag. Sci.* **57**, 1334–1349 (2011).
12. L. Eeckhoudt, H. Schlesinger, Putting risk in Its proper place. *Am. Econ. Rev.* **96**, 280–289 (2006).
13. S. T. Trautmann, G. van de Kuilen, Higher order risk attitudes: A review of experimental evidence. *Eur. Econ. Rev.* **103**, 108–124 (2018).
14. S. Olschewski, B. R. Newell, Y. Oberholzer, B. Scheibehenne, Valuation and estimation from experience. *J. Behav. Decis. Making* **34**, 729–741 (2021).
15. C. J. Burke, P. N. Tobler, Reward skewness coding in the insula independent of probability and loss. *J. Neurophysiol.* **106**, 2415–2422 (2011).
16. W. Genest, W. R. Stauffer, W. Schultz, Utility functions predict variance and skewness risk preferences in monkeys. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 8402–8407 (2016).
17. S. Shafir, A. Bechar, E. U. Weber, Cognition-mediated coevolution–context-dependent evaluations and sensitivity of pollinators to variability in nectar rewards. *Plant Syst. Evol.* **238**, 195–209 (2003).
18. N. Barberis, M. Huang, Stocks as lotteries: The implications of probability weighting for security prices. *Am. Econ. Rev.* **98**, 2066–2100 (2008).
19. O. K. Bayrak, J. D. Hey, Decisions under risk: Dispersion and skewness. *J. Risk Uncertain.* **61**, 1–24 (2020).
20. P. Bordalo, N. Gennaioli, A. Shleifer, Salience theory of choice under risk. *Q. J. Econ.* **127**, 1243–1285 (2012).
21. M. Dertwinkel-Kalt, M. Köster, Salience and skewness preferences. *J. Eur. Econ. Assoc.* **18**, 2057–2107 (2020).
22. R. Hertwig, I. Erev, The description-experience gap in risky choice. *Trends Cogn. Sci.* **13**, 517–523 (2009).
23. J. Tripp, G. D. A. Brown, Being paid relatively well most of the time: Negatively skewed payments are more satisfying. *Mem. Cognit.* **44**, 966–973 (2016).

24. E. U. Weber, S. Shafir, A.-R. Blais, Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychol. Rev.* **111**, 430 (2004).
25. G. Barron, I. Erev, Small feedback-based decisions and their limited correspondence to description-based decisions. *J. Behav. Decis. Making* **16**, 215–233 (2003).
26. A. Glöckner, B. E. Hilbig, F. Henninger, S. Fiedler, The reversed description-experience gap: Disentangling sources of presentation format effects in risky choice. *J. Exp. Psychol. Gen.* **145**, 486–508 (2016).
27. D. Kellen, T. Pachur, R. Hertwig, How (in)variant are subjective representations of described and experienced risk and rewards? *Cognition* **157**, 126–138 (2016).
28. C. K. Becker, E. Ert, S. T. Trautmann, G. van de Kuilen, Experiencing risk: Higher-order risk attitudes in description- and experience-based decisions. *J. Exp. Psychol. Learn. Mem. Cogn.* **47**, 727–746 (2021).
29. D. U. Wulff, M. Mergenthaler-Canseco, R. Hertwig, A meta-analytic review of two modes of learning and the description-experience gap. *Psychol. Bull.* **144**, 140–176 (2018).
30. C. Frydman, C. F. Camerer, The psychology and neuroscience of financial decision making. *Trends Cogn. Sci.* **20**, 661–675 (2016).
31. S. Olschewski, A. Luckman, A. Mason, E. A. Ludvig, E. Konstantinidis, The future of decisions from experience: Connecting real-world decision problems to cognitive processes. *Perspect. Psychol. Sci.*, 1–21 (2023).
32. S. Andersen, T. Hanspal, K. M. Nielsen, Once bitten, twice shy: The power of personal experiences in risk taking. *J. Fin. Econ.* **132**, 97–117 (2019).
33. J. J. Choi, D. Laibson, B. C. Madrian, A. Metrick, Reinforcement learning and savings behavior. *J. Fin.* **64**, 2515–2534 (2009).
34. U. Malmendier, S. Nagel, Depression babies: Do macroeconomic experiences affect risk taking? *Q. J. Econ.* **126**, 373–416 (2011).
35. M. Kaustia, S. Knüpfer, Do investors overweight personal experience? Evidence from IPO subscriptions. *J. Fin.* **63**, 2679–2702 (2008).
36. T. Lejarraga, J. K. Woike, R. Hertwig, Description and experience: How experimental investors learn about booms and busts affects their financial risk taking. *Cognition* **157**, 365–383 (2016).
37. S. Olschewski, L. Diao, J. Rieskamp, Reinforcement learning about asset variability and correlation in repeated portfolio decisions. *J. Behav. Exp. Fin.* **32**, 100559 (2021).
38. M. A. S. Bradbury, T. Hens, S. Zeisberger, Improving investment decisions with simulated experience. *Rev. Fin.* **19**, 1019–1052 (2015).
39. C. Kaufmann, M. Weber, E. Haisley, The role of experience sampling and graphical displays on one's investment risk appetite. *Manag. Sci.* **59**, 323–340 (2013).
40. T. Lohrenz, K. McCabe, C. F. Camerer, P. R. Montague, Neural signature of fictive learning signals in a sequential investment task. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 9493–9498 (2007).
41. M. S. Spektor, S. Gluth, L. Fontanesi, J. Rieskamp, How similarity between choice options affects decisions from experience: The accentuation-of-differences model. *Psychol. Rev.* **126**, 52–88 (2019).
42. K. Tsetsos, N. Chater, M. Usher, Salience driven value integration explains decision biases and preference reversal. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 9659–9664 (2012).
43. K. Tsetsos *et al.*, Economic irrationality is optimal during noisy decision making. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 3102–3107 (2016).
44. G. Loomes, R. Sugden, Regret theory: An alternative theory of rational choice under uncertainty. *Econ. J.* **92**, 805–824 (1982).
45. B. A. Mellers, A. Schwartz, K. Ho, I. Ritov, Decision affect theory: Emotional reactions to the outcomes of risky options. *Psychol. Sci.* **8**, 423–429 (1997).
46. W. M. Hayes, D. H. Wedell, Regret in experience-based decisions: The effects of expected value differences and mixed gains and losses. *Decision* **8**, 277–294 (2021).
47. D. Cohen, O. Plonsky, I. Erev, On the impact of experience on probability weighting in decisions under risk. *Decision* **7**, 153–162 (2020).
48. S. Shafir, T. Reich, E. Tsur, I. Erev, A. Lotem, Perceptual accuracy and conflicting effects of certainty on risk-taking behaviour. *Nature* **453**, 917–920 (2008).
49. S. Ebert, On skewed risks in economic models and experiments. *J. Econ. Behav. Organ.* **112**, 85–97 (2015).
50. I. Erev, A. E. Roth, Maximization, learning, and economic behavior. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 10818–10825 (2014).
51. I. Erev, E. Ert, O. Plonsky, D. Cohen, O. Cohen, From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychol. Rev.* **124**, 369–409 (2017).
52. A. Mason, C. R. Madan, N. Simonsen, M. L. Spetch, Biased confabulation in risky choice. *Cognition* **229**, 105245 (2022).
53. G. Schwarz, Estimating the dimension of a model. *Ann. Stat.* **6**, 461–464 (1978).
54. R. S. Sutton, A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, 2nd ed., 2018).
55. G. Molinaro, A. G. E. Collins, Intrinsic rewards explain context-sensitive valuation in reinforcement learning. *PLoS Biol.* **21**, e3002201 (2023).
56. W. M. Hayes, D. H. Wedell, Testing models of context-dependent outcome encoding in reinforcement learning. *Cognition* **230**, 105280 (2023).
57. L. Spiliopoulos, R. Hertwig, Variance, skewness and multiple outcomes in described and experienced prospects: Can one descriptive model capture it all? *J. Exp. Psychol. Gen.* **152**, 1188–1222 (2023).
58. A. Kraus, R. H. Litzenberger, Skewness preference and the valuation of risk assets. *J. Fin.* **31**, 1085–1100 (1976).
59. M. S. Spektor, D. U. Wulff, Myopia drives reckless behavior in response to over-taxation. *Judgm. Decis. Mak.* **16**, 114–130 (2021).
60. A. Parducci, Category judgment: A range-frequency model. *Psychol. Rev.* **72**, 407–418 (1965).
61. N. Stewart, N. Chater, G. D. A. Brown, Decision by sampling. *Cogn. Psychol.* **53**, 1–26 (2006).
62. I. Erev, E. Ert, O. Plonsky, Y. Roth, Contradictory deviations from maximization: Environment-specific biases, or reflections of basic properties of human learning? *Psychol. Rev.* **130**, 640–676 (2023).
63. I. Erev, O. Plonsky, Y. Roth, Complacency, panic, and the value of gentle rule enforcement in addressing pandemics. *Nat. Hum. Behav.* **4**, 1095–1097 (2020).
64. I. Erev *et al.*, Promoting healthy behavior through repeated deposit contracts: An intervention study. *J. Econ. Psychol.* **92**, 102548 (2022).
65. O. Yakobi, D. Cohen, E. Naveh, I. Erev, Reliance on small samples and the value of taxing reckless behaviors. *Judgm. Decis. Making* **15**, 266–281 (2020).
66. J. Baron, J. C. Hershey, Outcome bias in decision evaluation. *J. Pers. Soc. Psychol.* **54**, 569 (1988).
67. B. Spitzer, L. Waschke, C. Summerfield, Selective overweighting of larger magnitudes during noisy numerical comparison. *Nat. Hum. Behav.* **1**, 0145 (2017).
68. E. A. Ludvig, M. L. Spetch, Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS ONE* **6**, e20262 (2011).
69. F. Lieder, T. L. Griffiths, M. Hsu, Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychol. Rev.* **125**, 1–32 (2018).
70. Y. Vanunu, J. M. Hotaling, B. R. Newell, Elucidating the differential impact of extreme-outcomes in perceptual and preferential choice. *Cogn. Psychol.* **119**, 101274 (2020).
71. L. Weiss-Cohen, E. Konstantinidis, M. Speekenbrink, N. Harvey, Task complexity moderates the influence of descriptions in decisions from experience. *Cognition* **170**, 209–227 (2018).
72. T. Bali, D. Hirshleifer, L. Peng, Y. Tang, "Attention, social interaction, and investor attraction to lottery stocks" (Tech. Rep., National Bureau of Economic Research, Cambridge, MA, 2021).
73. E. Schulz, N. T. Franklin, S. J. Gershman, Finding structure in multi-armed bandits. *Cogn. Psychol.* **119**, 101261 (2020).
74. R. A. Rescorla, A. R. Wagner. "A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement" in *Classical Conditioning II*, A. H. Black, W. F. Prokasy, Eds. (Appleton-Century-Crofts, New York, NY, 1972), pp. 64–99.
75. K. Tsetsos, "The dynamics of selective integration during rapid experiential decisions" in *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*, G. Gunzelmann, A. Howes, T. Tenbrink, E. J. Davelaar, Eds. (Cognitive Science Society, 2017), pp. 1230–1235.
76. E.-J. Wagenmakers, A practical solution to the pervasive problems of p values. *Psychon. Bull. Rev.* **14**, 779–804 (2007).
77. M. S. Spektor, S. Olschewski, G. Le Mens, Frequent winners explain apparent skewness preferences in experience-based decisions–data and code. OSF. https://osf.io/aqjdz/. Deposited 8 February 2024.