Journal of Economic Behavior and Organization 221 (2024) 570–585

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Research Paper

Gambling in risk-taking contests: Experimental evidence ☆

Matthew Embrey [a,*], Christian Seel [b], J. Philipp Reiss [c]

[a] *University of Sussex, Department of Economics, Jubilee Building, Falmer, Brighton, BN1 9SL, UK*
[b] *Maastricht University, Department of Economics, PO Box 616, Maastricht, 6200 MD, Netherlands*
[c] *Karlsruhe Institute of Technology (KIT), Institute of Economics, Bluecherstr. 17, Karlsruhe, 76139, Germany*

ARTICLE INFO

JEL classification:
C72
C92
D81

Keywords:
Contests
Relative performance pay
Risk-taking behaviour
Laboratory experiment

ABSTRACT

This paper experimentally investigates risk taking in contest schemes by implementing a stopping task based on Seel and Strack (2013). In this stylized setting, managers with contest payoffs have an incentive to delay halting projects with a negative expectation, with the induced inefficiency being highest for a moderately negative drift. The experiment systematically varies the negative drift (between-subjects) and the payoff incentives (within-subject). We find evidence for risk taking in all our treatment conditions, with the non-monotonicity at least as pronounced as predicted. Contrary to the theoretical predictions, a similar pattern of behaviour persists even without contest incentives, suggesting contest incentives are not the only driver of risk-taking behaviour. Instead, observed behaviour violates expected utility maximization and is consistent with some intrinsic motivation for taking risk and myopic reasoning about the opponent. We explore the interplay between intrinsic incentives and the payment scheme and find that contest incentives might crowd out an intrinsic inclination to risk-taking.

1. Introduction

Relative performance schemes, such as contests and tournaments, are commonly used and important incentive schemes in many economic and social contexts. Such schemes can incentivise costly effort and select better-performing agents (Lazear and Rosen, 1981), while being simple to implement and credible even in cases where exact performance is not verifiable. Examples abound, including competition for bonuses or promotion, research and development races, status contests and fund-manager competition.

However, in many situations of interest, risky prospects to the contestants are a crucial factor. For instance, top managers decide which risky project to pursue and fund managers choose how to manage their portfolio (see, for example, Falkenstein, 1996; Chevalier and Ellison, 1999; Huang et al., 2011). They are often subject to contest incentives through bonus payments based on relative performance or job promotion opportunities (Kempf and Ruenzi, 2008). Moreover, funds compete for future cash inflow that is strongly correlated with past relative performance (Chevalier and Ellison, 1997), also leading to contest-type incentives.

☆ We wish to thank Dan Friedman, and Curtis Kephart and Alexander Wittmond from the LEEPS lab for permission to use and adapt the ConG software, and for their extensive support and assistance in developing it to implement the risk-taking stopping task. We greatly benefited from feedback on presentations in the ESA North American Conference in Dallas, the Behavioural-Experimental Social Science Workshop at the University of Essex, the PET conference in Strasbourg, the Behavioural Game Theory Workshop at the University of East Anglia, and seminars at the University of Nottingham, the University of Sussex, the University of Surrey, Bar Ilan University, University of Essex and Indiana University. Finally, we would like to thank Ferdinand Pieroth and Thorsten Rueger for excellent research assistantship.

* Corresponding author.
E-mail addresses: m.embrey@sussex.ac.uk (M. Embrey), c.seel@maastrichtuniversity.nl (C. Seel), philipp.reiss@kit.edu (J. Philipp Reiss).

https://doi.org/10.1016/j.jebo.2024.02.024

Available online 11 April 2024

As first explored in Hvide (2002), such contests are theoretically predicted to induce a high propensity to take risks. The subsequent theoretical literature, starting with Seel and Strack (2013), has identified two robust predictions about behaviour in a general class of contests in which agents can influence the risk: (i) risk-taking is pervasive and (ii) the potential losses for the principal are non-monotonic in the drift.[1]

This paper investigates experimentally these two predictions by implementing a novel stopping task based on Seel and Strack (2013): each subject privately observes a stochastic process over time—an independent realization of a random walk with drift—which could, for example, be thought of as representing the value of a fund manager's risky assets.[2] Each subject starts with the same initial value and is forced to stop if the value becomes zero (bankruptcy). The strategy of a subject is to specify when to stop the process and, under contest incentives, the subject who stops at the highest value wins a fixed prize (for example, a bonus for the fund manager).

In our implementation, we focus on the risk-taking incentives for self-regarding agents induced by different payment schemes.[3] As a result we do not include a principal (for example the stakeholders) as a passive third player who receives the stopped values. Under pure self-interest, the incentives for the agents in both models are equivalent, although observed behaviour in the experiment might change due to, for example, other-regarding preferences. The aim for this study was to control for heterogeneity in these homegrown preferences to focus on the role of contest incentives for self-regarding agents.

To test the theoretical predictions, we vary the drift parameters between-subjects, where the drift is minimally, moderately, or extremely negative. For each of these parameters, we additionally run a within-subject variation we call the individual-choice lottery treatment. In this treatment, theory predicts stopping immediately to be optimal for any expected utility maximizer. It provides a measure of a subject's inclination towards risk-taking, when facing the same stochastic process as in the contest but without the strategic incentive, as well as giving a direct contrast between lottery and contest payoff incentives. As such, the lottery treatment provides an "elimination design" (Niederle, 2015) without contest incentives to see whether the incentive scheme is the main driver of risk-taking as predicted by the theory.

We find evidence for risk-taking in all our contest treatment conditions as predicted by the theory. Furthermore, the non-monotonicity is at least as pronounced as predicted, with a moderately negative drift condition resulting in stopped values that are significantly smaller than those from both a minimally negative condition and an extremely negative condition. However, contrary to the benchmark predictions, we find similar aggregate patterns when we remove contest incentives. As such, our findings suggest that contest incentives are not the only cause of risk-taking behaviour in the experiment.

Further analysis of behaviour over periods reveals that, for the extreme negative drift, contestants are less likely to stop immediately under individual choice rather than contest incentives, which deviates in a systematic way from the theory predictions. With a moderate negative drift—the treatment condition that most exposes subjects to risk-taking incentives for strategic motives—we do not observe this crowding out of intrinsic motives effect. Instead, there is evidence subjects are more inclined to not stop immediately for strategic reasons. Observed behaviour could be ex-post rationalized by non-expected utility models such as cumulative prospect theory without the reduction axiom or a decreasing joy of gambling with a myopic reasoning about the opponent, but more future work is needed to isolate the channels.

In terms of methodology, our paper introduces a novel experimental task to study risk-taking contests. To see this, note that a stopping strategy (e.g. "stop after 10 seconds" or "stop whenever the process reaches value 5 or 17") induces a probability distribution over the stopped value of the process. This distribution contains all the relevant strategic information as there is no dynamic interaction between the players. Two alternative implementations are possible. First, a "distribution-builder" approach (Sharpe et al., 2000), where subjects move probability mass back and forth subject to feasibility, i.e., there exists a stopping time which implements the desired distribution. Second, a selection task from a pre-selected list of feasible distributions such as the binary distributions induced by threshold strategies. In our setting, our implementation is closer to the theory and it seems more interactive and less tedious than the other options. Yet, it would be interesting to compare behaviour under different implementations in future work.

There is a large set of applications suitable to such experiments. A simple example would be an all-pay auction with a linear cost function, where subjects first choose their effort cost, and then can either choose from a pre-selected list of distributions with this effort cost or implement the stopping task without drift. But even beyond contest games, there are many possible applications—for example, on redistributive politics (Myerson, 1993), public good provision under different electoral incentives (Lizzeri and Persico, 2001), and general lotto games (Hart, 2008).

Related literature

The seminal theoretical paper on relative performance schemes is Lazear and Rosen (1981), who argue that the optimal contest induces the first-best effort level under risk neutrality and might outperform other simple payment schemes such as a piece rate if

[1] The related literature adapts the model of Seel and Strack (2013) to more general stochastic processes (Feng and Hobson, 2015), asymmetric bankruptcy constraints (Seel, 2015), incomplete information about the endowment (Feng and Hobson, 2016), flow costs of research (Seel and Strack, 2016), multiple prizes with an arbitrary structure (Fang and Noe, 2016; Strack, 2016; Fang et al., 2020; Fang and Noe, 2022), partial observability and a Black-Scholes model rather than a simple stopping problem (Strack, 2016).

[2] The theoretical model of Seel and Strack (2013) uses a continuous time approach with a Brownian motion. We discretize the state space and choose a random walk that changes its value 4 times per second to approximate the process in discrete time. In line with the literature, we argue that due to the high frequency of observations, the behaviour of subjects should not differ systematically from the continuous-time setting.

[3] Note that the incentives for risk-taking in this setting are not dependent on the risk attitude of the decision maker; see Remark 1 for details.

agents are risk averse. Hvide (2002) challenges these findings by allowing contestants to determine the variance of their performance measure at no cost; in equilibrium, agents choose zero effort, but an infinite level of variance. Nieken and Sliwka (2010) implemented a simplified binary version (high or low risk) of the model by Hvide (2002) in the laboratory. Their focus lies on head-starts by one opponent. In line with the theory predictions, they find that laggards tend to take higher risks. These models and their implementation are static, although most applications in which agents choose risk in contests are inherently dynamic. The model of Seel and Strack (2013) addresses these concerns by analyzing a contest in continuous time in which agents get dynamic feedback and can adapt their risk choices over time.

Discretisations of continuous-time frameworks have been used to study behaviour in the lab in wide variety of strategic settings, including network formation (Berninghaus et al., 2006), hawk-dove games (Oprea et al., 2011; Berninghaus et al., 2012), the prisoner's dilemma game (Friedman and Oprea, 2012; Bigoni et al., 2015), minimum-effort game (Deck and Nikiforakis, 2012), rock-paper-scissors (Cason et al., 2014) and Hotelling competition (Kephart and Friedman, 2015). Probably the closest setting to ours is the preemption game experiment of Anderson et al. (2010), where subjects also observe a random walk with short time steps. Yet, in contrast to our setting, all subjects in a group see the same random walk. Additionally, there is a positive probability that it stops in each period, which results in no payoff if no subject stopped the process before. Otherwise, the first to stop the process receives the current (common) valuation minus a private cost.

Random walk models have also been used in individual choice experiments examining topics including the timing of investment decisions (Oprea et al., 2009), earning withdrawals (Oprea, 2014), asset liquidation (Magnani, 2015), costly price adjustment (Magnani et al., 2016), and optimal stopping with regret (Strack and Viefers, 2021).

Most of the experimental contest literature focuses on effort as the crucial variable of interest. Notable exceptions include Dijk et al. (2014) and Kirchler et al. (2018) that investigate the role of rank incentives and social comparison on risk-taking behaviour in a repeated portfolio-choice settings. The repeated setting generates interim laggards and leaders, and Dijk et al. (2014) find that, with a student subject population, laggards invest in positively skewed assets and leaders in negatively skewed assets. Kirchler et al. (2018) find that financial professionals increase risk-taking with both tournament and rank incentives, while students only increase risk taking with tournament incentives. Contrary to our study, risk taking increases the expected return in these settings, and the multi-period decision-making produces interim feedback on relative performance.

Eriksen and Kvaløy (2017) consider a setting in which the optimal strategy is always to take no risk, and find that risk-taking is nonetheless increasing in the competitiveness of a tournament. In our setting, risk-taking is an integral part of predicted behaviour with contest incentives but does not increase expected returns. Our focus is on how extrinsically-motivated risk-taking interacts with economic conditions; although, intrinsic motivations for risk-taking do play an important role in understanding our results.

2. Theoretical background

In this section, we provide a brief review of the theoretical results from Seel and Strack (2013) that are relevant for our experimental implementation. In their setup, each of two agents $i = 1, 2$ controls a project whose value is determined by a stochastic process $X^i = (X^i_t)_{t\in\mathbb{R}_+}$. The stochastic processes are governed by the law of motion $X^i_t = x_0 + \mu t + \sigma^2 B^i_t$, where $x_0 > 0$ is the common starting value of each process at time $t = 0$, $\mu < 0$ is the common drift parameter, $\sigma^2 > 0$ is the common variance parameter and B^1_t and B^2_t are independent Brownian motions. Note that since $\mu < 0$, each process decreases in expectation.

At every point in time, each agent privately observes the value of her own project. A strategy of the agent specifies a plan when and at which values she irreversibly stops her process. This decision can depend on the previous realisation of the agent's process, but since the agent obtains no information about her rival, it does not depend on the realisation or stopping decision of the rival. If the value of the project becomes zero (bankruptcy), the agent is forced to stop. The game ends when both agents have stopped their processes. The agent whose stopped value is higher wins a prize. In case of a tie, each agent wins the prize with probability 50%. Each agent maximizes her expected payoff, which is the probability of winning the contest times the prize.

While each player's decision problem is dynamic, the game itself is static since no new information about the rival arrives over time. Thus, Seel and Strack (2013) focus on the Nash equilibria of the game. To state their main characterisation result, denote the probability distribution over the stopped process which is induced by a stopping strategy of player i by F^i; i.e., $F^i(x)$ is the probability that player i stops her process below or at the value x. The following result sums up Propositions 1-3 in Seel and Strack (2013).[4]

Proposition 1. *In any Nash equilibrium, both players choose strategies which induce the distribution*

$$F^1(x) = F^2(x) = F(x) = \min\left\{ \frac{1}{2} \frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{\exp(\frac{-2\mu x_0}{\sigma^2}) - 1}, 1 \right\}. \tag{1}$$

[4] Seel and Strack (2013) also allow for positive values of μ, while we already restrict attention to $\mu < 0$ in the model description. The distributions are unique for $n = 2$ (see Proposition 3 in Seel and Strack, 2013).

Thus, there is a unique prediction for the distribution of the stopped value. From the distribution F, Seel and Strack (2013) compute the expected value of each player's stopped process and obtain the following main comparative statics result (Proposition 7 in Seel and Strack, 2013).

Proposition 2. *The expected value of the stopped processes in equilibrium is first falling and then rising in the drift (U-Shape) and attains its minimum at a moderately negative level of the drift.*

To get an intuition for why contestants would not always stop immediately in equilibrium, consider a strategy profile in which both players stop at the starting value and win with probability $\frac{1}{2}$. Consider a deviation such that one contestant continues until the stopped value decreases by a large amount or increases by a small amount, where the latter occurs with a higher probability. In the case it increases (which happens with a probability over $\frac{1}{2}$), she wins the game since her rival stopped at a lower value. Thus, we have constructed a profitable deviation, meaning a profile in which both players stop immediately for sure is not a Nash equilibrium. The equilibrium distribution is constructed such that no matter which strategy the rival chooses, her winning probability is at most $\frac{1}{2}$.

We close the section with two properties which will be useful for the experimental implementation and robustness analysis.

Remark 1 *(Independence of the size of the prize and the risk attitude)*. In an expected utility setting, maximizing the expected payoff (prize times winning probability) yields the same result as maximizing the winning probability. In particular, this result requires the reduction axiom, i.e., compound lotteries can be reduced to simple lotteries. In this framework, the theoretical prediction is independent of the size of the prize. Moreover, maximizing the winning probability does not depend on a player's risk attitude.

Remark 2 *(Skewed stopped value distribution in equilibrium)*. The distribution function in Eq. (1) has a strictly increasing density on its support. Thus, the mode is at the upper bound of the support and the mean is below the median (x_0). There are several measures of skewness which all agree that the distribution is left-skewed. As one example, consider the widely used Pearson's median skewness defined as $3(\text{mean} - \text{median})/\sigma$ which is below 0 as the mean is below the median.

3. Experimental design

The baseline implementation of the contest game uses the following setup: Two agents compete for a prize of 150 ECU.[5] At every point in time an agent privately observes the value of her own project. While the stochastic process that determines this value is the same for both agents, each agent has a separate and independent realization from this process. The game ends after 90 seconds. If the process has not been stopped until that time, the stopped value is equal to the final value of the process.

The implemented process is a random walk which starts at a value of 15. Every quarter of a second, the value of the process either moves up by one with probability p or moves down by one with probability $1 - p$. Subjects were explicitly given these probabilities in the instructions, along with sample realizations. This implementation aims to find a balance between two counteracting potential problems. Making the discretization too fine—or equivalently starting value too high—can lead to a time limit of 90 seconds, at which a player is forced to stop in the experiment, being a binding restriction. On the other hand, making the discretization too coarse can lead to the incentives for not stopping immediately being eliminated, especially for more negative drift parameters.[6] In the limit as the grid size and time converge to zero, the random walk converges to the Brownian motion setting analyzed in Seel and Strack (2013); see Appendix C of the supplementary materials for the mapping from the discrete-time process parameters to the continuous-time process parameters. As the related literature (e.g., Anderson et al., 2010, Pettit et al., 2014, and Oprea, 2014), we argue that the frequency of changes in the process is rapid enough that subjects perceive it to be near continuous-time, and thus their behaviour is not qualitatively different from the continuous time setting.

The only action available to an agent is to choose when to stop the project. This could be done by pressing a "stop now" button or by setting upper and/or lower thresholds for the value of the project. The thresholds would trigger automatic stopping if the value was greater than or equal to the threshold, in the case of the upper threshold, or less than or equal to the threshold, in the case of the lower threshold. We use the thresholds to give subjects a convenient way to stop the first time the process hits a certain value, i.e., not having to react on the spot in that case. See Fig. 1 for a screenshot of the computer interface for this stopping task. The agent is forced to stop if the value of the project hits zero. The agent whose stopped value is the highest is awarded the prize. In case of a tie, each agent wins with probability 50%.

3.1. Treatment design

The experiment has two primary goals. The first is to test the prediction that contest incentives do indeed induce decision-makers to delay stopping projects that lose money in expectation. The second is to test the prediction that the expected losses from this

[5] Throughout, all payoffs are in experimental currency units, denoted ECU, which are converted at a rate of 0.01 Euros per ECU.

[6] In contrast to the benchmark theory, there is an exogenous upper bound on the length of the contest game. In the supplementary materials, we provide several arguments why the time bound should not constrain predicted behaviour in our implementation. Our parameters are chosen such that waiting for at least one uptick or bankruptcy is always a best response against stopping immediately.

Fig. 1. Screenshot: Active Phase.

Table 1
Summary of Between-Subject Treatment Conditions.

Treatment	Abbreviation	Pr(up)	Expected Stopped Value
Minimal Negative	Min-ve	0.49625	14.54
Moderate Negative (baseline)	Mod-ve	0.47	13.71
Extreme Negative	Ext-ve	0.3375	14.53

Notes: The value of Pr(up) gives the probability of an increase in value of the process (random walk with $X_0 = 15$, time interval jump size $\Delta X = 1$ and time interval $\Delta t = 0.25$). The underlying Brownian motion (continuous time) has $X_0 = 15, \sigma = 2$ and $\mu \in \{-0.03, -0.24, -1.3\}$ for Min-ve, Mod-ve, and Ext-ve, respectively.

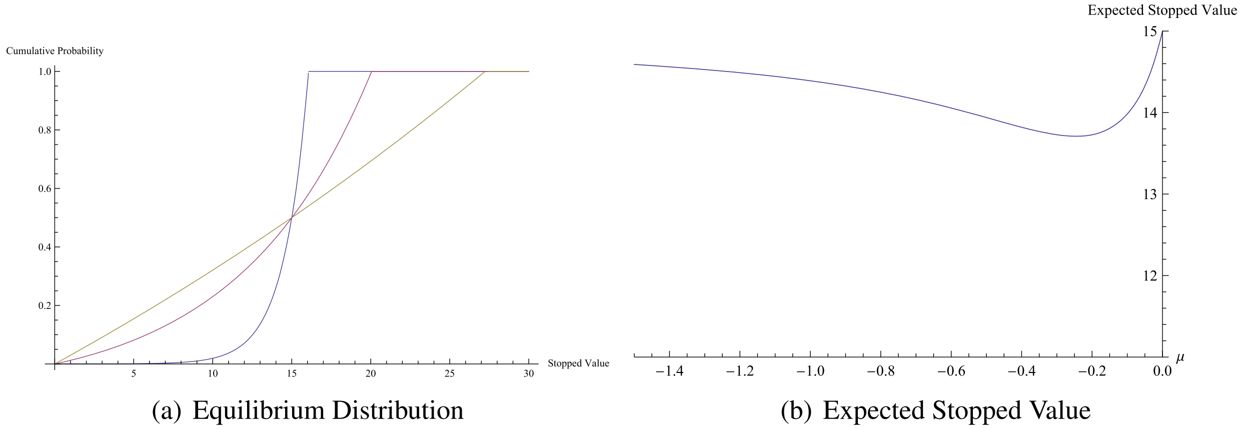

Fig. 2. Equilibrium distribution (panel a) and expected stopped value (panel b) for the parameters ($x_0 = 15, \sigma^2 = 2$) depending on the drift μ; the implemented drift parameters for (a) are $\mu = -1.3$ (blue graph), $\mu = -0.24$ (red graph), $\mu = -0.03$ (yellow graph), while (b) covers $\mu \in [-1.5, 0]$. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

risk-taking are not monotonic in the fundamentals. The project fundamental is the drift parameter that governs the expected loss of continuing the project. We consider three between-subject treatments with different drift values, a minimal-negative, extreme-negative and moderate-negative baseline as shown in Table 1. The corresponding equilibrium distributions from the model are displayed in the left panel of Fig. 2, while the stopped values are displayed in the right panel.

For every treatment, a within-subject variation is also included where the winning probability is determined according to a linear function of the stopped value of their project. This variation switches off the strategic element of the contest, reducing the environment to an individual choice setting. It provides a measure of a subject's inclination towards gambling, when facing the same stochastic process as in the contest but without the strategic incentive, as well as giving a direct contrast between lottery and contest payoff incentives.

In the lottery variation, the subject controls a project whose value follows exactly the same random process as in the contest game they experience. As before, their only decision is when to stop the project. Now, rather than being used to determine the winner of a tournament with another subject, the stopped value of the project is used to determine the probability the subject will win a prize worth 150 ECU. The probability of winning is given by the following linear rule

$$\text{Probability of winning 150 ECU prize} = \min\left\{\frac{X}{30}, 1\right\},$$

where X is the stopped value of a subject's project. This payoff scheme was chosen since it best mimics the scheme obtained in the contest: should both subjects choose to stop their projects immediately then there is a 50-50 chance of winning the prize in both the contest and the individual choice setting. Furthermore, as in the contest setting, the maximisation problem in the individual choice setting is independent of the size of the prize and the risk attitude.

Since the process decreases in expectation and the winning probability is linear in the stopped value, the theoretical prediction is that subjects stop their process immediately. The following result formalises this intuition; the proof is relegated to the appendix.

Proposition 3. *To maximize the probability of winning the prize, the individual must stop immediately at the starting value in the individual choice setting.*

3.2. Predictions

We have three main hypotheses based on the theory. The calculation of the point predictions for the implemented parameters is explained in Appendix C of the supplementary materials.

Hypothesis 1. (a) In the contest, players do not always stop immediately.
(b) The average time until the process is stopped increases in the drift.
(c) The average time until the process is stopped equals 0.72 seconds for Ext-ve, 10.75 seconds for Mod-ve, and 30.66 seconds for Min-ve.

Hypothesis 2. (a) The average stopped value is non-monotone in the drift, first falling and then rising (U-Shape).
(b) The average stopped value equals 14.53 for Ext-ve, 13.71 for Mod-ve, and 14.54 for Min-ve.

Hypothesis 3. The process is stopped immediately in the individual choice setting.

3.3. Procedures

The experiment was conducted in the BEElab at Maastricht University between December 2014 and April 2015. 240 students were recruited using ORSEE (Greiner, 2004, 2015) and participated in one of the three treatments. During each session, up to three matching groups of 8 subjects were run in parallel. Sessions lasted 90 minutes on average. For each treatment variation, ten matching groups were run, split evenly over two order combinations (see below). Each matching group comprised of eight participants.

A session consisted of three parts. The first part was an instruction part during which subjects were given details of the structure of the session and the environment in which they were to make decisions (see Appendix D of the supplementary materials for an example of the instructions); subjects were also given experience with the stochastic process that underlies the contest game, as well as the interface for stopping the stochastic process. After the first part of the instructions were read out aloud,[7] subjects were shown graphs of the complete path for ten randomly drawn example realisations—these were block randomised so that every subject in the same matching group was shown the same set of examples. The examples were static in the sense that they showed the complete realised path, rather than having them drawn in real-time. After the ten realised graphs, subjects were asked a series of comprehension questions to make sure that they understood the basic features of the stochastic process—such as, the probability that the process may go up, that each jump up or down is independent of any previous realisations, and the consequences of the process hitting zero. The graphs and comprehension questions, as well as a basic demographics questionnaire at the end of the session, were implemented in zTree (Fischbacher, 2007).

Once everyone had correctly answered the comprehension questions, the session switched to ConG (Pettit et al., 2014), the main software interface in which the dynamic stopping task was implemented (see Appendix E of the supplementary materials for additional screenshots of the computer interface). To give subjects real-time experience of the process, as well as experience using the various controls for stopping it, the instructional part of each session concluded with five trial periods. During the trial periods,

[7] Subjects were given the complete set of instructions at the beginning of the session. However, the software was paused at the end of the first and second parts in order to read out aloud the instructions that were specific for the upcoming part. Without pausing, the software would have otherwise transitioned through the different parts (trial-contest-lottery or trial-lottery-contest, depending on the task-order variation) without any warning for subjects.

Table 2
Summary of treatments.

Treatment	Task Order	Number of Matching Groups	Subjects	Earnings (Euro) Min.	Max.	Avg.
Min -ve	Contest-Lottery	5	40	10.50	25.50	17.32
Min -ve	Lottery-Contest	5	40	12.00	27.00	17.51
Mod -ve	Contest-Lottery	5	40	10.50	22.50	16.65
Mod -ve	Lottery-Contest	5	40	10.50	24.00	16.46
Ext -ve	Contest-Lottery	5	40	10.50	27.00	17.77
Ext -ve	Lottery-Contest	5	40	10.50	25.50	17.66

Table 3
Summary of the Outcomes in the Contest.

Treatment	Percent Stopped Time > 0		Stopped Time (sec) Average (95% C.I.)		Predicted	Stopped Value (ECU) Average (95% C.I.)		Predicted
Min -ve	95.5	***	23.74	(18.12, 29.36)	30.66	14.69	(13.60, 15.79)	14.54
Mod -ve	82.0	***	8.78	(6.78, 10.79)	10.75	12.86	(12.35, 13.37)	13.71
Ext -ve	26.2	***	0.25	(0.16, 0.34)	0.72	14.70	(14.55, 14.85)	14.53

Notes: Data from the last five contest periods. *** denotes significantly different from zero with p-value < 0.001. Tests and confidence intervals use standard errors clustered at the matching group level.

the random process was run five times independently for each subject using the main software interface. The trial periods were not paid, and subjects were neither matched into pairs nor given any feedback on other participants' realisations or stopped values.

In the baseline (Contest-Lottery) order combination, part two consisted of ten periods of the contest game, played under the same treatment conditions. At the beginning of a period the computer interface randomly matched participants within a matching group into pairs. The period begins with an empty graph showing just the axes and the interface buttons. The graph starts at 15 and remains there for the first 15 seconds, so that subjects have an opportunity to prepare for the start of the period (the warm-up phase)—at this point they can set their upper and lower thresholds in preparation for the start of the random fluctuation, or they can press the "stop now" to stop the process at its start value. After the warm up time, the graph starts to randomly fluctuate according to an independent realization of the random walk described above. This fluctuation continues for 90 seconds, during which subjects can stop the process either automatically using the threshold controls or by pressing the "stop now" button. Subjects continue to see the realisation of the random walk even after they have fixed their stopped value for that period. Once the 90 seconds are over, subjects are shown the stopped value of the person they were matched with and the outcome of the period. This feedback phase lasts 15 seconds before the next period begins.

During part three, subjects played ten periods of the individual choice version of the contest game, referred to here as the lottery task. The stochastic process and interface were exactly as in the contest task, only the payment and feedback were adjusted to reflect the fact that this was an individual (rather than a paired) task with a lottery to determine the final payment. After part three was completed, participants were paid in cash according to the amount of ECUs they accumulated during part two and three plus a small show-up fee of 3 Euro. Paying out every period of play reduces variance of payments for the subjects.[8] For the reverse order combination (Lottery-Contest), the lotteries were in part two and the contests in part three. Table 2 gives a summary of the sessions.

4. Main experimental results

This section provides the main experimental results, in particular the tests of the hypotheses stated in Section 3.2. Before these explicit tests, Fig. 3 gives an overview of behaviour in the last five periods of the contest and lottery parts of the experiment. The top graphs show the observed distribution of the stopped time for the three drift parameters. As can be seen, there is a clear ranking across treatments, with the whole distribution of stopped times getting shorter as the drift parameter becomes more negative. Contrary to the theoretical predictions, aggregate behaviour is strikingly similar in the contest and the lottery treatment.

The lower graphs give the observed distribution for the stopped value, using the same data range and breakdown. Again, there is a clear separation across treatments, with the distribution of stopped values becoming less dispersed as the drift parameter becomes more negative. For the contest, this pattern qualitatively resembles that predicted, although there are some quantitative differences to the distribution function strictly predicted by the benchmark theory: the observed upper bound of the support is larger, particularly for the Min-ve and Mod-ve treatments, and there appear to be larger mass points at bankruptcy (stopped value equal zero) and at the start value (stopped value equal 15).

Table 3 gives average values in the last five contest periods for the three outcome variables that will be used to explicitly test the main hypotheses: an indicator for the stopped time being strictly greater than zero, the stopped time and the stopped value. Similar

[8] This payment scheme does not disturb incentives since subjects' stopping decisions only affect the probability of winning in every period and no hedging is possible.

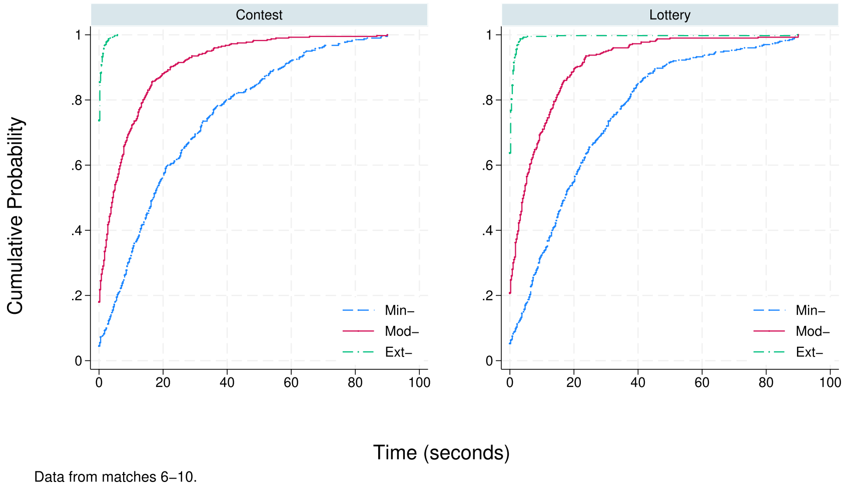

(a) Stopped Time

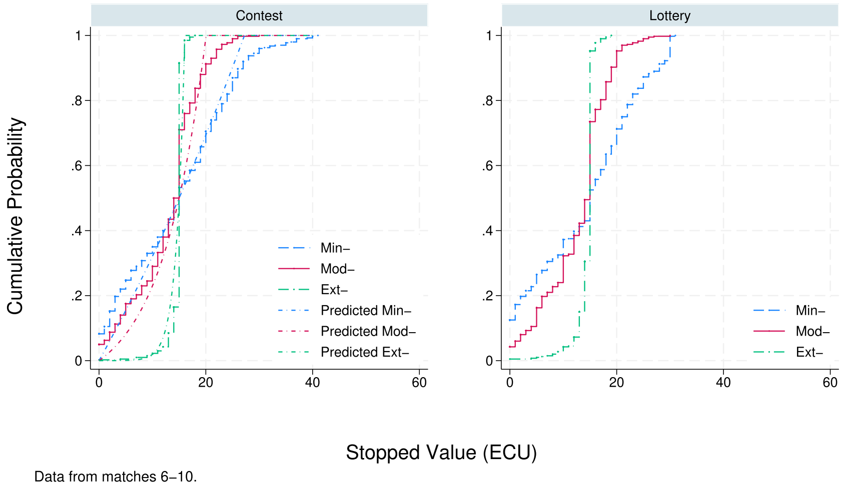

(b) Stopped Value

Fig. 3. The Observed Distribution of Stopped Times and Stopped Values.

conclusions are drawn when considering a "between-subjects" test that uses periods 11-15 from just the contest-lottery sessions for the contest outcomes, and periods 11-15 from just the lottery-contest sessions for the lottery outcomes.[9]

The averages for the percent of observations with a stopped time strictly greater than zero show that a significant proportion of subjects do not stop immediately for all drift parameters, in line with Hypothesis 1A. The statistical significance for this is established using a random-effects probit regression on the probability of not stopping immediately with a complete set of treatment dummies as independent variables. To allow for the potential correlation between observations caused by the matching scheme in the contest part of the experiment, the standard errors are clustered at the matching-group level.[10] Positive point probabilities of stopping immediately are consistent with the theory once we discretize the state space. As to be expected, they increase as the drift gets more negative.

[9] The main text focuses on the last five periods of the contest and lottery tasks. See Tables F.1 and F.2 of the supplementary materials for a comparison of these outcomes for the first five periods and over all periods. See Table F.3 for the details of the between-subjects comparison.

[10] This gives ten clusters per treatment, five for each of the two task orders. Throughout the results, statistical significance will generally be established using an analogous approach—that is, by constructing an appropriate regression and using clustered-robust standard errors, corrected to allow arbitrary correlation within a matching-group, as well as a subject-specific random effect. This is not the only approach to addressing the issue of potential session-effects in experimental data. A common alternative is to take matching-group averages—see Fréchette (2011) for a discussion of session-effects. Robustness checks using non-parametric tests on matching-group averages are reported in Tables B.7 and B.8 of the appendix. They lead to the same conclusions.

Table 4
Summary of the Outcomes in the Lottery.

	Percent Stopped		Stopped Time (sec)			Stopped Value (ECU)		
Treatment	Time > 0		Average (95% C.I.)		Predicted	Average (95% C.I.)		Predicted
Min -ve	94.8	***	22.62	(20.03, 25.22)	0.00	14.24	(13.51, 14.97)	15.00
Mod -ve	79.2	***	8.49	(6.94, 10.03)	0.00	12.73	(12.11, 13.34)	15.00
Ext -ve	36.2	***	0.61	(0.19, 1.03)	0.00	14.36	(14.17, 14.54)	15.00

Notes: Data from the last five lottery periods. *** denotes significantly different from zero with p-value < 0.001. Tests and confidence intervals use standard errors clustered at the matching group level.

The stopped time column shows that the time spent before stopping reduces as the drift parameter gets more negative, with the average stopped time in the Min-ve treatment significantly longer than in the Mod-ve treatment (p-value < 0.001), which in turn is significantly longer than in the Ext-ve treatment (p-value < 0.001)—see Table B.8 for further details and robustness checks. These results are again in line with Hypothesis 1A. For the point predictions from Hypothesis 1A, subjects spend significantly less time before stopping in both the Min-ve and Ext-ve treatments than theoretically predicted, while the average stopped time for the Mod-ve treatment is just within the 95% confidence interval—see Table F.4 of the supplementary materials for explicit tests of the point predictions.

Result 1. *Subjects do not always stop the process immediately in the contest. The average time before stopping in the contest reduces as the drift becomes more negative.*

The final column of Table 3 gives the average stopped value across treatments in the last five contest periods. As predicted in Hypothesis 2A, the stopped value is non-monotonic in the drift parameter with the average in the Mod-ve treatment being significantly smaller than in both the Min-ve and Ext-ve treatments, which have statistically similar averages—see Table B.8 for further details and robustness checks. Indeed the U-shape prediction is more pronounced than predicted as the average in the Mod-ve treatment is significantly smaller than the point prediction, while the other treatments are in line with their point prediction—see Table F.4 of the supplementary materials for explicit tests of the point predictions.

Result 2. *The average stopped value is non-monotonic in the drift parameter. In particular, the stopped value in the contest is lowest for the moderately negative drift parameter. The observed non-monotonicity is more pronounced than predicted.*

Table 4 repeats the analysis of Table 3 for the last five lottery periods. While the evidence from the contest is mostly in line with the associated hypotheses, this is no longer the case in the lottery treatments. There is a significant amount of not stopping immediately in all treatments, contrary to Hypothesis 3 (see also Table B.7). This behaviour responds to the treatment condition, with the average stopped time reducing as the drift parameter gets more negative. Furthermore, the non-monotonicity in average stopped values is also observed in the lottery task (see also Table B.8 for these last two points). All point predictions for the lottery are rejected (see also Table F.4 of the supplementary materials).

Result 3. *Subjects do not always stop the process immediately in the lottery. The average time before stopping in the lottery reduces as the drift becomes more negative. Furthermore, the stopped value in the lottery is also non-monotonic in the drift parameter.*

The within-subject variation can also be used to examine differences in subjects' behaviour under the contest incentives compared to the lottery incentives. Table 5 reports this analysis, using data from the last five periods of the contest and lottery parts.[11] The difference columns report the average subject-level difference in their contest average minus their lottery average for each of three main outcome variables. The analysis suggests some differences in the within-subject contrast across the different drift parameters that will be explored in further detail in the subsequent section. First, there are no significant differences between subjects' average behaviour in the contest and average behaviour in the lottery for Min-ve treatment. At the other end of the spectrum, there are significantly *higher* stopped values in the contest rather than the lottery for the Ext-ve treatment, with subjects stopping sooner (in particular for the lottery-contest task order). The pattern is different again for Mod-ve treatment, with significantly less stopping immediately in the contest than the lottery under the contest-lottery task order, although less time spent before stopping in the contest in the lottery-contest task order.

[11] See Figures F.1–F.3 of the supplementary materials for scatter plots of the subject averages across the different parts of the experiment. This analysis can also be done using the panel structure of the data set, rather than averaging at the subject level, and controlling for the period trends. This approach is taken in Section 5 and the analogous within-session comparisons are reported in Table F.6 of the supplementary materials. This table also reports the between-session contrast that compares behaviour in the contest by those that played the contest first with behaviour in the lottery by those that played the lottery first. In both cases, the analysis leads to the same qualitative results: noisy behaviour in the Min-ve treatment and more stopping immediately in the contest than the lottery in the Ext-ve, while this pattern is not seen in the Mod-ve.

Table 5
Within-Subject Differences in Outcomes.

	Percent Stopped Time > 0		Stopped Time (sec)		Stopped Value (ECU)	
Treatment/Task Order	Difference	p-value	Difference	p-value	Difference	p-value
Min-ve						
Contest-Lottery	0.0	1.000	3.73	0.272	0.42	0.732
Lottery-Contest	1.5	0.636	-1.50	0.510	0.49	0.726
All	0.7	0.654	1.12	0.600	0.45	0.614
Mod-ve						
Contest-Lottery	12.5	0.001	2.40	0.114	-0.49	0.324
Lottery-Contest	-7.0	0.167	-1.81	0.040	0.75	0.351
All	2.7	0.518	0.30	0.781	0.13	0.794
Ext-ve						
Contest-Lottery	5.5	0.399	-0.50	0.212	0.28	0.073
Lottery-Contest	-25.5	0.001	-0.22	0.012	0.41	0.098
All	-10.0	0.139	-0.36	0.076	0.34	0.016

Notes: Data from the last five periods of the contest phase and the last five periods of the lottery phase. The difference columns report the average subject-level differences between the contest average and the lottery average. The p-value columns report the result of a two-sided test against the null hypothesis that on average subject-level differences between the contest average and the lottery average is equal to zero. The test is based on a linear regression on treatment indicator variables with standard errors clustered at the matching group level.

5. Further analysis

This section analyses subjects' behaviour in more detail; specifically, considering the observed dynamics over periods, the heterogeneity in subject risk-taking and the determinants of risk-taking choices under contest incentives. The focus here is on the main results for the extensive margin, the percentage with stopping time greater than zero. The supplementary materials contain further details on the analysis, as well as the analogous analysis for the overall stopped time and stopped-value distribution. Not stopping immediately under contest incentives, but stopping immediately under lottery incentives, is the most basic prediction of the benchmark, purely rational model of behaviour in this environment. While there are also predictions for overall stopped time and the stopped-value distribution, they are necessarily noisier given that they are more dependent on the realisation of the random process. Furthermore, some important predictions for the stopped-value distribution involve higher-order moments of the distribution than the mean (in particular, the skewness), something we return to after discussing potential explanations both within and outside the expected utility framework.

5.1. Further analysis of behaviour

Fig. 4(a) shows the development of the likelihood of not stopping immediately across periods. The graph is split both by drift parameter (different panels) and by task order (different lines in a panel). To quantify these observations, Table 6 reports the results of a random-effects probit regression for each drift parameter using the payoff structure indicator (contest versus a baseline of lottery) and order indicator (lottery-contest versus a baseline of contest-lottery) and the interaction term, along with a period trend variable. A simple linear trend across parts is used given the similar dynamics of the propensity to not stop immediately across the contest and lottery periods; see Table F.7 of the supplementary materials for a robustness check on the trend specification. Fig. 4(b) shows the predicted behaviour from this regression model.[12] To complete the picture, Fig. 5 gives the average behaviour across periods for (overall) stopped time and stopped value.[13]

In the Min-ve treatment, there is little difference in the propensity to not stop immediately—which is in general quite high—between the incentive conditions and order combinations, or over time. At the other end of the spectrum, there is a strong trend to stop immediately more often with experience in the Ext-ve treatment. Here, for a given level of experience in terms of periods, subjects are more likely to stop immediately under the contest payoff structure than under the lottery payoff structure, especially

[12] In addition to a drift-parameter specific trend, the regressions allow for four types of test, or contrast, regarding the incentive scheme for each drift-parameter value. Two of these are "within-subject" in that they concern behaviour across parts of the same session, while two of these are "between-subject" in that they concern behaviour in the same part across sessions. The first within-subject contrast is the average effect of contest incentives in the standard contest-lottery order, and is given by the contest coefficient. The second within-subject contrast is the average effect of contest incentives in the reverse lottery-contest order, and is given by the sum of the contest and contest × pay order coefficients. The first between-subject contrast is the average effect of contest incentives in part 2 compared to lottery incentives in part 2, and is given by the *difference* between the contest and pay order coefficients. The other between-subject contrast is the analogous effect of contest incentives in part 3, and is given by the sum of the pay order, contest and contest × pay order coefficients. Given subjects in part 2 had not experienced another incentive scheme, the first of these gives the "purer" between-subjects comparison. For reference, Table F.6 of the supplementary materials gives the complete set of point predictions and tests for the trend and these four incentive scheme contrasts for all three outcome variables.

[13] Analogous linear random-effects regressions were run for these; see Section F.1 of the supplementary materials.

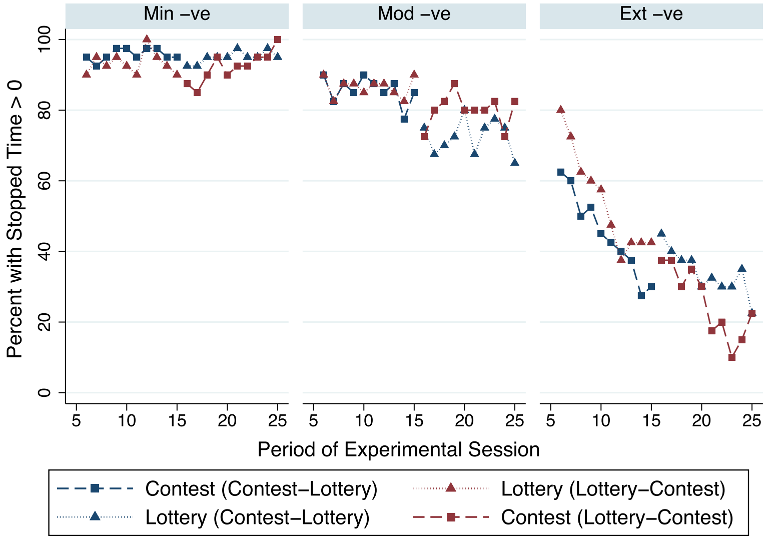

(a) Observed

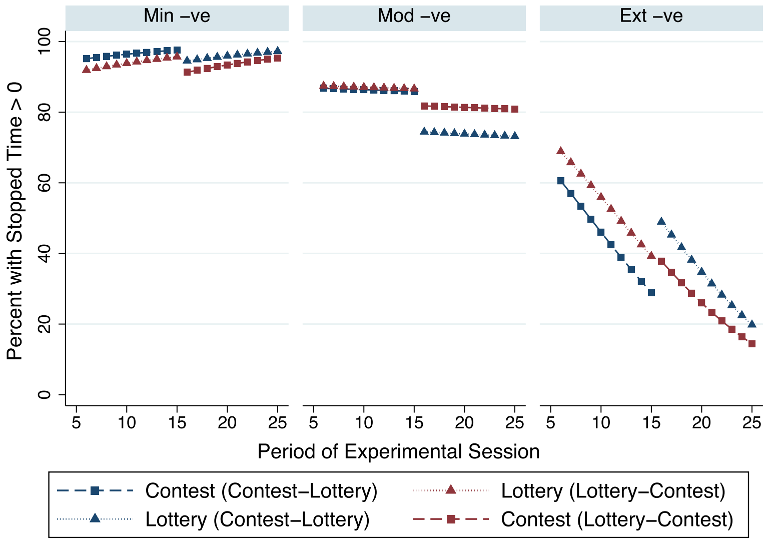

(b) Predicted

Fig. 4. Behaviour across Periods by Task Order: Percentage with Stopped Time > 0.

for the contest-lottery order where there is a significant increase in the propensity to not stop immediately following the switch to lottery payoffs. This effect is consistent with competition under contest payoffs driving out intrinsic motives for risk-taking.

In the Mod-ve treatment, there is no strong trend to stop immediately more often over periods. Furthermore, there is no suggestion that this increases under the lottery payoffs as seen in the Ext-ve condition. Instead there is a increase in stopping immediately following the switch of regime in the contest-lottery order; under the lottery-contest order, this propensity remains roughly the same before and after the switch.

To investigate the heterogeneity in subjects' general inclination towards not stopping immediately under the contest incentive scheme, we use the observed behaviour from the lottery periods to construct an individual measure (the proportion of periods a subject's stopped time is strictly greater than zero in the lottery) of a subject's preference for risk-taking. This individual measure and an indicator for whether the subject's stopped time was strictly greater than zero in the first match of the contest are used to explicitly model subject risk-taking types. In all treatments these individual components play a significant role in the probability of not stopping immediately in a contest period (see Model 1 of Table F.11 of the supplementary materials). This specification is then extended to include explanatory variables for whether a subject's opponent from the previous match had a stopped value equal to 15

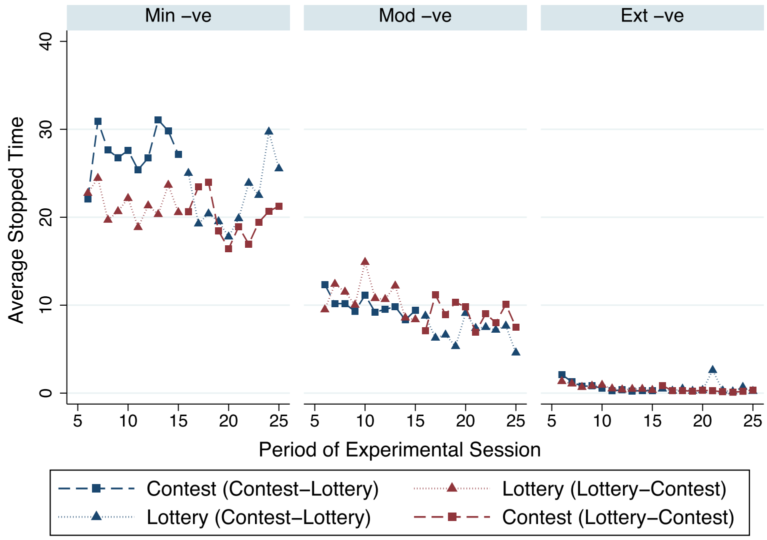

(a) Stopped Time

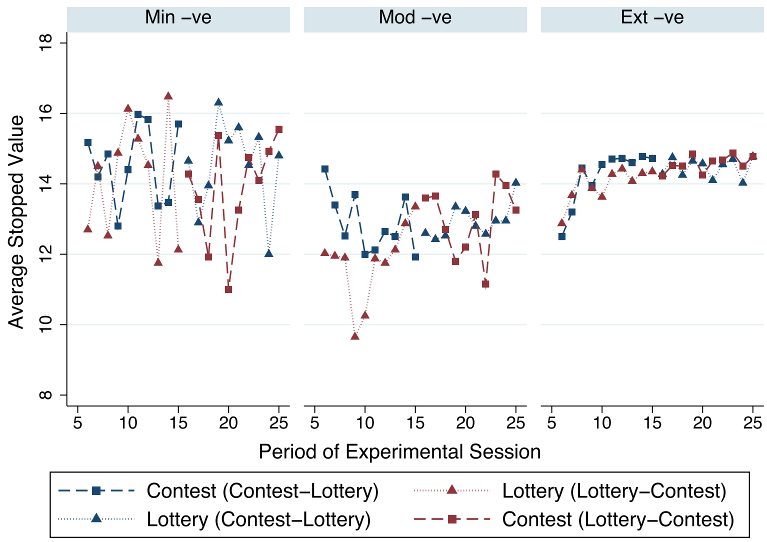

(b) Stopped Value

Fig. 5. Observed Behaviour across Periods by Task Order.

Table 6
Probability Stopped Time > 0 across Periods Regressions.

	Min-ve		Mod-ve		Ext-ve	
Contest	0.02	(0.188)	0.08***	(0.004)	-0.26***	(0.001)
Lottery-Contest Order	0.01	(0.623)	0.10**	(0.042)	-0.15**	(0.027)
Contest × Lottery-Contest Order	-0.04	(0.154)	-0.12***	(0.009)	0.28*	(0.088)
Period	0.00	(0.181)	-0.00	(0.760)	-0.04***	(0.000)

Notes: Random-effects probit models. Coefficients report the marginal effect on the dependent variable. p-value of significance test in parentheses using standard errors clustered at the matching group level. *** 1%, ** 5%, * 10%. The period variable is normalised to start at 1 for the first for-payment period (which is period 6 of a session, the first period of part 2 after the purely trial periods of part 1) and finishes at 20 (which is period 25 of a session, the last period of part 3).

or strictly greater than 15. While these feedback variables have the expected sign in all treatments, it is only in the Mod-ve condition that it has a significant bearing on subjects' choice to not stop immediately (see Model 2 of Table F.11). The supplementary materials provide a more detailed description of the evidence on the heterogeneity in subject behaviour and on the possible determinants of not stopping immediately under contest incentives (see Section F.1). The following results sum up the findings from this section.

Result 4. *In all treatments, there is evidence for individual heterogeneity in the propensity to not stop immediately consistent with part of the risk-taking being due to individual preferences.*

Result 5. *For both lottery and contest incentives there is a significant trend to stop immediately more often with experience in the Ext-ve treatment, while there is no such trend in the Min-ve and Mod-ve treatments. In the Ext-ve treatment, contest payoffs appear to crowd out intrinsic motives for risk-taking. The same is not true in the Mod-ve treatment, with evidence that subjects are more inclined to stop immediately under lottery incentives.*

Result 6. *In the Mod-ve treatment, subjects respond significantly to the revealed outcome of their match under contest incentives—being more likely to stop immediately if their opponent's stopped value was equal to the start value, and being less likely to stop immediately if it is above the stopped value.*

5.2. Explanations both within and outside the expected utility framework

The individual choice treatment features the most stark difference from the theoretical predictions. We consider different possible explanations outside the classical expected utility framework. One important model in the class is (cumulative) prospect theory (Kahneman and Tversky, 1979), where agents use a weighting function to evaluate probabilities. In related models of (casino) gambling such as Barberis (2012) and Ebert and Strack (2015) the use of cumulative prospect theory leads to different predictions than expected utility, as agents often derive a high utility from constructing skewed gambles.

In our framework, a similar approach does not change the main predictions: agents should stop immediately in the lottery task. To understand why, recall that any stopping strategy induces a binary probability distribution over winning or not winning the prize. The probability of winning the prize is highest if an agent stops immediately, i.e., this option first-order stochastically dominates every other strategy. Note that this result holds true regardless whether the agent is sophisticated or naive. Yet, the main ingredient to obtain our prediction is the reduction axiom: agents reduce compound lotteries to simple lotteries; see Segal (1990) for a detailed account. If, unlike the above mentioned literature, we allow agents to apply the probability weighting to each step, we could rationalize strategies which do not stop immediately under cumulative prospect theory.

In Section C of the supplementary materials, we consider several other models including regret, joy of gambling and myopic reasoning about the opponent. A model with a decreasing joy of gambling and myopic reasoning about the opponent is able to bridge a lot of gaps between the theoretical predictions and observed behaviour. Consistent with the data, but unlike other explanations, this model is able to obtain risk-taking in the lottery and to obtain skewed distributions in the contest. We emphasize that the model was chosen based on ex-post reasoning on the observed data. Thus, one should be careful not to over interpret the results. Furthermore, the approach of including the intrinsic motive for gambling as a function of the amount of time spent exposed to the random fluctuations before stopping would also incorporate other possible intrinsic motivations from risk-taking in the experiment that would be observationally equivalent, such as a simple activity bias.[14]

A striking prediction in Seel and Strack (2013) is the convexity and left-skew of the equilibrium stopped-value distribution. The behavioural implication of this prediction is to distort the stopped value distribution in the hope of winning by a small amount while risking losing by a large amount (the "win-small-lose-big" strategy). Intrinsic motives for risk-taking that are tied simply to the amount of time spent exposed to the random fluctuation would not have this feature (see Remark 3 of Section C of the supplementary materials). In Section F.1 of the supplementary materials, we find evidence that the stopped value distribution for the Mod-ve treatment is negatively skewed under contest incentives, but not under lottery incentives.

6. Conclusion

This paper presents a theory-based experimental investigation of behaviour in risk-taking contests using a novel laboratory stopping task. Despite its dynamic nature, the experimental framework presents subjects with a relatively simple and intuitive decision-making problem that implements a choice among a large set of probability distributions, while maintaining a tight connection to the underlying theory. Predicted behaviour in our experiment is independent of the risk attitude of the participants, which reduces potential confounds and allows us to compensate every single period of the experiment without distorting incentives. From a methodological point of view, the implementation of choices over probability distributions as a stopping decision might also pave the way for future experimental research in other areas such as distortionary taxation or general lotto games.

The two main theoretical predictions for the contest— risk-taking and expected losses which are non-monotone in the drift—are observed in the experiment. In line with the theory, we find evidence for risk-taking in all our treatment conditions. Furthermore, the non-monotonicity is at least as problematic as predicted, with a moderately negative drift condition resulting in stopped

[14] See, for example, Lugovskyy et al. (2010), who consider the role of active participation in over-bidding in an experimental all-pay auction setting.

values that are significantly smaller than those from both a minimally negative condition and an extremely negative condition.

Contrary to the benchmark theoretical predictions, we also find this aggregate pattern of behaviour even when we remove the contest incentives and instead determine the winning probability as a linear function of the stopped value. Further analysis suggests, along with a great deal of heterogeneity, that many subjects display behaviour consistent with some intrinsic motivation for taking-risk in the stopping task. However, the intrinsic motive and the contest motive for risk-taking appear to reinforce the non-monotonicity in different ways: Under the extreme negative condition, contest incentives appear to crowd out the intrinsic motive for risk taking, while under the moderate negative condition contest payoffs increase the propensity to take risks. Such behaviour is consistent with decreasing joy of gambling in combination with myopic reasoning about the stopping decision.

Thus, we have observed that there might be risk-taking under the contest scheme, but the dependence on the compensation scheme is more sensitive to the parameters than predicted. Given the importance of compensation schemes for top managers and fund managers, behaviour induced by these schemes in risky environments needs to be well-understood. This paper has uncovered an interesting behavioural nuance—contest incentives might crowd out an intrinsic inclination for risk-taking—but it is only a starting point towards improving behavioural predictions for risky environments by combining empirical evidence with theoretical models.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Appendix A. Additional proofs

Proof of Proposition 3. Let us first show that stopping immediately is an optimal strategy:

Recall that if the process is stopped at a value X_t, the winning probability is $\min\left\{\frac{X_t}{30}, 1\right\} = \frac{1}{30}\min\left\{X_t, 30\right\}$. Thus, the player chooses a stopping time τ to solve $\max_\tau \frac{1}{30}\min\left\{X_\tau, 30\right\}$.

Note that $X = (X_t)_{t\in\mathbb{R}_+}$ is a supermartingale, i.e., since $\mu < 0$, the process is decreasing in expectation. Thus, by Doob's optional sampling theorem (see, e.g., Peskir and Shiryaev, 2006, p.60), for any bounded stopping time τ, $\mathbb{E}(X_\tau) \leq X_0$. Thus, for any bounded stopping time τ, we obtain $\mathbb{E}(\frac{1}{30}\min\left\{X_\tau, 30\right\}) \leq \frac{1}{30}\mathbb{E}(X_\tau) \leq \frac{1}{30}X_0 = \frac{1}{2}$. Hence, it is an optimal strategy to stop immediately, since this results in a winning probability of $\frac{1}{2}$. We do not spell out explicitly that stopping at $t = 0$ is *strictly* better than any continuation strategy $\tau > 0$, but refer the reader to Theorem 1.6 in Bismut and Skalli (1977) who prove the result by characterizing the optimal stopping region for supermartingales.

Appendix B. Robustness checks for main hypothesis tests

Table B.7
Probability Stopped Time > 0: Non-Parametric Tests.

Treatment	Percent Stopped Time > 0	Regression based p-value	K-S p-value based on: All data	Subject	Matching group
Hypothesis 1a: Probability stopped time > 0 in the contest					
Min -ve	95.50	0.000	0.000	0.000	0.000
Mod -ve	82.00	0.000	0.000	0.000	0.000
Ext -ve	26.25	0.000	0.000	0.000	0.000
Hypothesis 3: Probability stopped time > 0 in the lottery					
Min -ve	94.75	0.000	0.000	0.000	0.000
Mod -ve	79.25	0.000	0.000	0.000	0.000
Ext -ve	36.25	0.000	0.000	0.000	0.000

Notes: The regression-based *p*-values are based on a linear random-effects regression (with clustering at the matching group level) of the stopped time or stopped value on a set of treatment indicators. The K-S *p*-value is based on a Kolmogorov-Smirnov test between the degenerate distribution that puts all weight on stopping immediately and the observed distribution using either all the data, subject averages or matching-group averages. Data from the last five periods of the contest and the lottery parts, respectively.

Table B.8
Treatment Comparisons in Hypotheses 1–3: Non-Parametric Tests.

Comparison	Regression based p-value	Rank sum p-value	K-S p-value based on: All data	Subject	Matching group
Hypothesis 1b: Stopped time in the contest					
Min -ve vs Mod -ve	0.000	0.000	0.000	0.000	0.000
Min -ve vs Ext -ve	0.000	0.000	0.000	0.000	0.000
Mod -ve vs Ext -ve	0.000	0.000	0.000	0.000	0.000
Hypothesis 2a: Stopped value in the contest					
Min -ve vs Mod -ve	0.003	0.008	0.000	0.000	0.052
Min -ve vs Ext -ve	0.989	0.289	0.000	0.000	0.418
Mod -ve vs Ext -ve	0.000	0.000	0.000	0.000	0.000
Hypothesis 3: Stopped time in the lottery					
Min -ve vs Mod -ve	0.000	0.000	0.000	0.000	0.000
Min -ve vs Ext -ve	0.000	0.000	0.000	0.000	0.000
Mod -ve vs Ext -ve	0.000	0.000	0.000	0.000	0.000
Hypothesis 3: Stopped value in the lottery					
Min -ve vs Mod -ve	0.002	0.010	0.000	0.003	0.052
Min -ve vs Ext -ve	0.759	0.623	0.000	0.000	0.418
Mod -ve vs Ext -ve	0.000	0.002	0.000	0.000	0.002

Notes: The regression-based *p*-values are based on a linear random-effects regression (with clustering at the matching group level) of the stopped time or stopped value on a set of treatment indicators. The rank-sum *p*-value is based on the ranks of matching-group averages. The K-S *p*-value is based on a Kolmogorov-Smirnov between the observed distributions of either all the data, subject averages or matching group averages. All tests use data from the last five periods of the contest and the lottery parts, respectively.

Appendix C. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jebo.2024.02.024.

References

Anderson, S.T., Friedman, D., Oprea, R., 2010. Preemption games: theory and experiments. Am. Econ. Rev. 100, 1778–1803.
Barberis, N., 2012. A model of casino gambling. Manag. Sci. 58 (1), 35–51.
Berninghaus, S.K., Ehrhart, K.-M., Ott, M., 2006. A network experiment in continuous time: the influence of link costs. Exp. Econ. 9, 237–251.
Berninghaus, S.K., Ehrhart, K.-M., Ott, M., 2012. Forward-looking behaviour in hawk-dove game in endogenous networks: experimental evidence. Games Econ. Behav. 75, 35–52.
Bigoni, M., Casari, M., Skrzypacz, A., Spagnolo, G., 2015. Time horizon and cooperation in continuous time. Econometrica 83, 587–616.
Bismut, J.-M., Skalli, B., 1977. Temps d'arrêt optimal, théorie générale des processus et processus de markov. Probab. Theory Relat. Fields 39, 301–313.
Cason, T.N., Friedman, D., Hopkins, E., 2014. Cycles and instability in a rock-paper-scissors population game: a continuous time experiment. Rev. Econ. Stud. 81, 112–136.
Chevalier, J., Ellison, G., 1997. Risk taking by mutual funds as a response to incentives. J. Polit. Econ. 105, 1167–1200.
Chevalier, J., Ellison, G., 1999. Carrier concerns of mutual fund managers. Q. J. Econ. 114, 389–432.
Deck, C., Nikiforakis, N., 2012. Perfect and imperfect real-time monitoring in a minimum-effort game. Exp. Econ. 15, 71–88.
Dijk, O., Holmen, M., Kirchler, M., 2014. Rank matters–the impact of social competition on portfolio choice. Eur. Econ. Rev. 66, 97–110.
Ebert, S., Strack, P., 2015. Until the bitter end: on prospect theory in a dynamic context. Am. Econ. Rev. 105, 1618–1633.
Eriksen, K.W., Kvaløy, O., 2017. No guts, no glory: an experiment on excessive risk-taking. Rev. Finance 21 (3), 1327–1351.
Falkenstein, E.G., 1996. Preference for stock characteristics as revealed by mutual fund portfolio holdings. J. Finance 51, 111–135.
Fang, D., Noe, T., 2016. Skewing the odds: Taking risks for rank-based rewards. Working Paper.
Fang, D., Noe, T., 2022. Less competition, more meritocracy? J. Labor Econ. 40 (3), 669–701.
Fang, D., Noe, T., Strack, P., 2020. Turning up the heat: the discouraging effect of competition in contests. J. Polit. Econ. 128 (5), 1940–1975.
Feng, H., Hobson, D., 2015. Gambling in contests modelled with diffusions. Decis. Econ. Finance 38 (1), 21–37.
Feng, H., Hobson, D., 2016. Gambling in contests with random initial law. Ann. Appl. Probab. 26, 186–215.
Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.
Fréchette, G., 2011. Session effects in the laboratory. Exp. Econ. 15, 485–498.
Friedman, D., Oprea, R., 2012. A continuous dilemma. Am. Econ. Rev. 102, 337–363.
Greiner, B., 2004. The online recruitment system ORSEE 2.0 - a guide for the organization of experiments in economics. Working Paper.
Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. J. Econ. Sci. Assoc. 1 (1), 114–125.
Hart, S., 2008. Discrete colonel blotto and general lotto games. Int. J. Game Theory 36, 441–460.
Huang, J., Sialm, C., Zhang, H., 2011. Risk shifting and mutual fund performance. Rev. Financ. Stud. 24, 2575–2616.
Hvide, H., 2002. Tournament rewards and risk taking. J. Labor Econ. 20, 877–898.
Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica 47, 263–292.
Kempf, A., Ruenzi, S., 2008. Tournaments in mutual-fund families. Rev. Financ. Stud. 21, 1013–1036.
Kephart, C., Friedman, D., 2015. Hotelling revisits the lab: equilibration in continuous and discrete time. J. Eur. Sci. Assoc. 1, 132–145.
Kirchler, M., Lindner, F., Weitzel, U., 2018. Rankings and risk-taking in the finance industry. J. Finance 73, 2271–2302.
Lazear, E.P., Rosen, S., 1981. Rank-order tournaments as optimal labor contracts. J. Polit. Econ. 89, 841–864.
Lizzeri, A., Persico, N., 2001. The provision of public good under alternative electoral incentives. Am. Econ. Rev. 91, 225–239.

Lugovskyy, V., Puzzello, D., Tucker, S., 2010. An experimental investigation of overdissipation in the all pay auction. Eur. Econ. Rev. 54 (8), 974–997.
Magnani, J., 2015. Testing for the disposition effect on optimal stopping decisions. Am. Econ. Rev. Pap. Proc. 105 (5), 371–375.
Magnani, J., Gorry, A., Oprea, R., 2016. Time and state dependence in an Ss decision experiment. Am. Econ. J. Macroecon. 8, 283–310.
Myerson, R.B., 1993. Incentives to cultivate favored minorities under alternative electoral systems. Am. Polit. Sci. Rev. 87, 856–869.
Niederle, M., 2015. Intelligent design: the relationship between economic theory and experiments: treatment-driven experiments. In: Fréchette, G.R., Schotter, A. (Eds.), Handbook of Experimental Economic Methodology. Oxford University Press, NY.
Nieken, P., Sliwka, D., 2010. Risk-taking tournaments - theory and experimental evidence. J. Econ. Psychol. 31, 254–268.
Oprea, R., 2014. Survival versus profit maximization in a dynamic stochastic experiment. Econometrica 82, 2225–2255.
Oprea, R., Friedman, D., Anderson, S.T., 2009. Learning to wait: a laboratory experiment. Rev. Econ. Stud. 76, 1103–1124.
Oprea, R., Henwood, K., Friedman, D., 2011. Separating the hawks from the doves: evidence form continuous time laboratory games. J. Econ. Theory 146, 2206–2225.
Peskir, G., Shiryaev, A., 2006. Optimal Stopping and Free-Boundary Problems. Birkhäuser Verlag, Basel, Boston, Berlin.
Pettit, J., Friedman, D., Kephart, C., Oprea, R., 2014. Software for continuous game experiments. Exp. Econ. 17, 631–648.
Seel, C., 2015. Gambling in contests with heterogeneous loss constraints. Econ. Lett. 136, 154–157.
Seel, C., Strack, P., 2013. Gambling in contests. J. Econ. Theory 148, 2033–2048.
Seel, C., Strack, P., 2016. Continuous time contests with private information. Math. Oper. Res. 41 (3), 1093–1107.
Segal, U., 1990. Two-stage lotteries without the reduction axiom. Econometrica 58, 349–377.
Sharpe, W.F., Goldstein, D.G., Blythe, P.W., 2000. The distribution builder: a tool for inferring investor preferences, preprint.
Strack, P., 2016. Risk-taking in contests - the impact of fund-manager compensation on investor welfare. Working Paper.
Strack, P., Viefers, P., 2021. Too proud to stop: regret in dynamic decisions. J. Eur. Econ. Assoc. 19, 165–199
.