Gambling in Risk-Taking Contests: Experimental Evidence

Citation for published version (APA):

Embrey, M., Seel, C., & Reiss, J. P. (2020). *Gambling in Risk-Taking Contests: Experimental Evidence*. Maastricht University, Graduate School of Business and Economics. GSBE Research Memoranda No. 025 <https://doi.org/10.26481/umagsb.2020025>

Document status and date:

Published: 22/09/2020

DOI:

[10.26481/umagsb.2020025](https://doi.org/10.26481/umagsb.2020025)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Matthew Embrey, Christian Seel,
J. Philipp Reiss

**Gambling in Risk-Taking
Contests: Experimental Evidence**

RM/20/025

ISSN: 2666-8807

GSBE

Maastricht University School of Business and Economics
Graduate School of Business and Economics

P.O. Box 616
NL- 6200 MD Maastricht
The Netherlands

Gambling in Risk-Taking Contests: Experimental Evidence

Matthew Embrey*

University of Sussex

m.embrey@sussex.ac.uk

Christian Seel†

Maastricht University

c.seel@maastrichtuniversity.nl

J. Philipp Reiss§

KIT

philipp.reiss@kit.edu

August 11, 2020

Abstract

This paper experimentally investigates excessive risk taking in contest schemes by implementing a novel 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 excessive risk taking in all our treatment conditions, with the non-monotonicity at least as problematic as predicted. Contrary to the theoretical predictions, this aggregate pattern of behaviour is seen even without contest incentives. Further analysis suggests that many subjects display behaviour consistent with some intrinsic motivation for taking risk. This intrinsic motive and the strategic motive for excessive risk taking reinforce the non-monotonicity. The experiment uncovers a behavioural nuance where contest incentives crowd out an intrinsic inclination to gamble.

JEL classification: C72, C92, D81

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

*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 benefitted 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 and the University of Surrey. Finally, we would like to thank Ferdinand Pieroth and Thorsten Rueger for excellent research assistantship.

†Department of Economics, Jubilee Building, Falmer, United Kingdom.

‡Department of Economics, PO Box 616, 6200 MD Maastricht, The Netherlands.

§Karlsruhe Institute of Technology (KIT), Institute of Economics, Bluecherstr. 17, 76139 Karlsruhe, Germany.

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 effort is not necessarily the predominant choice variable for the contestants, with risk often being a crucial factor. This is particularly the case for higher-level managerial and fund management decisions, with the latter receiving a great deal of attention in recent literature. 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, Sialm, and Zhang, 2011). Nonetheless, 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), leading to contest-type incentives.

In such settings contest incentives are liable to induce excessive risk taking as first explored in a theoretical model by Hvide (2002). The subsequent theoretical literature has identified two robust predictions about behaviour in a general class of contests in which agents can influence the risk: (i) excessive risk-taking is pervasive (for example, Fang and Noe, 2016; Strack, 2016; Fang, Noe, and Strack, 2020) and (ii) the potential losses for the principal are highest if additional risk comes at a moderate expected loss (for example, Seel and Strack, 2013).

This paper investigates experimentally these two predictions by implementing a novel experimental stopping task: each subject privately observes a stochastic process over time—a random walk with drift, which could, for example, be thought of as representing the value of a fund manager’s risky assets. Each subject starts with the same initial value and is forced to stop if the value becomes zero (i.e. 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 promotion for the fund manager). In all treatment conditions the drift parameter is negative, meaning that the

efficient choice is stopping immediately, while excessive risk-taking is delaying in the hope of a sequence of positive realisations.

We find evidence for excessive 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 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 pay subjects linearly in their 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 excessive risk-taking appear to reinforce the non-monotonicity in subtly different ways: Under the extreme negative condition, contest incentives appear to crowd out the intrinsic motive for excessive-risk taking, while under the moderate negative condition contest payoffs increase the propensity to take risks.

The experimental task implements the gambling-in-contests model of Seel and Strack (2013). They derive the Nash equilibrium outcome of the contest, and show that agents take excessive risks in equilibrium. In particular, due to the relative performance incentives, they do not even stop the process immediately if it has a negative expectation. In this case, the expected value of the stopped process is quasi-convex in the drift, i.e., highest losses occur for moderately negative expectations. While the model of Seel and Strack (2013) is deliberately simple and highly stylized, the recent literature has shown that many qualitative predictions, such as the excessive risk-taking, extend if one allows for more general stochastic processes (Feng and Hobson, 2015), asymmetric bankruptcy constraints (Seel, 2015), incomplete information about the endowment (Feng and Hobson, 2016a; Fang and Noe, 2018), flow costs of research (Seel and Strack, 2016), multiple prizes with an arbitrary structure (Fang and Noe, 2016; Strack, 2016), partial observability and a Black-Scholes model rather than a simple stopping problem (Strack, 2016).

To test these predictions, we choose the setting of Seel and Strack (2013) which provides a relatively simple framework to test the two effects. Yet, the predictions are theoretically robust to the extensions mentioned in the previous paragraph. Moreover, this setting is particularly suited for an experimental implementation since predicted be-

haviour is independent of the risk attitude, a potentially confounding variable in many other experiments. In our setting, contestants only influence their probability of winning the prize. Thus, the theory only requires that each contestant prefers to increase his probability of winning the prize.

To implement the setting of Seel and Strack (2013) in the laboratory, we need a discrete and finite version of the contest game. Thus, we discretize the state space and approximate the Brownian motion by a random walk. The random walk changes its value 4 times per second. 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. We run three different treatments, where we systematically vary the drift parameter. Our baseline treatment implements the moderate negative drift parameter for which the theory predicts highest expected losses in equilibrium. The other two treatments implement an extreme negative and a minimal negative value of the drift. For each of these treatments, we run an additional within-subject individual choice treatment which switches off the contest incentive. In this treatment, the probability of the winning prize is linear in the stopped value.

Based on the theory predictions, we test three main hypotheses. First, in the contest, subjects should not stop immediately (Hypothesis 1). Second, expected losses are maximal for the moderately negative drift and about equal for the other two values (Hypothesis 2). Third, subjects should stop immediately in the within-subject individual choice treatment (Hypothesis 3). For the first two hypotheses, we also test the point predictions derived from the theory.

The implementation worked as follows. Subjects participated in 25 rounds of play. The first five rounds were non-incentivized, individual training rounds, which gave subjects the opportunity to familiarize themselves more with the random process and the computer interface. Then the contest with two randomly rematched contestants per group and the individual choice treatments were played for 10 rounds each. At the end of the experiment, subjects received 1,50 Euro for each round in which they won the prize and a show-up fee of 3 Euro. Note that paying every round of play does not disturb incentives in this experiment, since subjects' stopping decisions only affect the probability of winning in every round.

The main results of the experiment are: in line with the theory prediction, subjects

do not always stop immediately in the contest. Moreover, the expected value of the stopped process in the contest is quasi-convex in the drift, even more pronounced than the theory predicts. Thus, the experiment provides evidence in favour of Hypothesis 1 and Hypothesis 2. However, contrary to Hypothesis 3, there is also substantial evidence of gambling in the individual choice treatment, suggesting intrinsic motives for gambling as well as strategic ones.

Further analysis of behaviour over rounds reveals that, for the extreme negative drift, contestants gamble more (in terms of extensive margin) under individual rather than contest incentives, which deviates in a systematic way from the theory predictions. With a moderate negative drift—which is the treatment condition that most exposes subjects to the inefficient incentives to gamble for strategic motives—we do not observe this crowding out of intrinsic motives effect. Instead, there is evidence subjects are more inclined to gamble for strategic reasons. We also document quite some heterogeneity in subjects’ behaviour, both in terms of their general inclination towards gambling, as well their response to the change in incentive structure.

We consider two sets of behavioural theories that might help organize the data. The first theory is based on different models of regret over the failure to stop at a better time if there was a winning strategy in hindsight. Indeed, the computer interface suggests this, since subjects see the entire realisation of the stopping process irrespective of when they chose to stop. However, the predictions from these regret-based theories are not consistent with the observed pattern of behaviour across treatments. In particular, such regret motivations should only reinforce the zero-gambling-in-the-lottery prediction of the standard theory. The second theory is based on “joy of gambling”; that is, an additional procedural utility for not stopping the process. A simple model with constant joy of gambling does not help in understanding the data. A model of decreasing joy of gambling combined with myopic reasoning, however, is able to accommodate the experimental results for both the lottery and the individual choice treatment.

From a methodological point of view, this paper introduces a way to implement choices among a large set of probability distributions in the laboratory by a stopping decision. This method does not artificially reduce the choice problem by pre-selecting a list of feasible distributions. Furthermore, differing from an implementation with dragging mass back and forth between certain values, our method seems more natural and it can be used

if the expectation of the random variable differs for different feasible distributions. As such, our method might inspire future experimental research within and beyond contest games—for example, on redistributive politics (Myerson, 1993), public good provision under different electoral incentives (Lizzeri and Persico, 2001), and general lotto games (Hart, 2008).

After a more detailed review of the related literature, the paper proceeds as follows. Section 2 summarizes the model and theoretical predictions of Seel and Strack (2013) as far as they are relevant for the experimental design. Section 3 presents the experimental design, the procedures and primary hypotheses. The main experimental results, in particular with regard to the primary hypotheses, are reported in Section 4, while Section 5 provides more detailed analysis of behaviour. We discuss additional behavioural explanations based on regret and joy of gambling in Section 6. Section 7 concludes. All proofs, additional theoretical background material and supplementary figures and tables are relegated to an online appendix.

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 agents are risk averse (see, for example, Charness and Kuhn (2011) for a review of tournament experiments, in particular those using the static framework). 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. While the model by Hvide (2002) nicely illustrates that contests might induce excessive risk-taking, the assumption of costless choice of variance and the predictions are rather extreme. 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.

However, note that these models and the implementation thereof 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 con-

tinuous time in which agents get dynamic feedback and can adapt their risk choices over time.

A continuous-time framework has been used to study behaviour in the lab in wide variety of strategic settings, including network formation (Berninghaus, Ehrhart, and Ott, 2006), hawk-dove games (Oprea, Henwood, and Friedman, 2011; Berninghaus, Ehrhart, and Ott, 2012), the prisoner’s dilemma game (Friedman and Oprea, 2012; Bigoni, Casari, Skrzypacz, and Spagnolo, 2015), minimum-effort game (Deck and Nikiforakis, 2012), rock-paper-scissors (Cason, Friedman, and Hopkins, 2014) and Hotelling competition (Kephart and Friedman, 2015). Probably the closest setting to ours is the preemption game of Anderson, Friedman, and Oprea (2010). In their experiment, subjects observe a random walk with short time intervals, which is very similar to the one we use. However, 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, Friedman, and Anderson, 2009), earning withdrawals (Oprea, 2014), costly price adjustment (Magnani, Gorry, and Oprea, 2016), and optimal stopping with regret (Strack and Viefers, 2019).

Finally, there is a large recent experimental literature on contests, in particular lottery contests, all-pay auctions and rank-order tournaments—see Dechenaux, Kovenock, and Sheremeta (2015), for a survey. Most of this literature focuses on effort as the crucial variable of interest. A notable exception is Kirchler, Lindner, and Weitzel (2018) who analyze binary risk choices in a multi-period setting with (non-monetary) rank incentives, tournament incentives, and without incentives. They 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 Kirchler, Lindner, and Weitzel (2018). Furthermore, in their paper, the total size of the payoff depends on the actions of the participants, even in the tournament incentive treatment.

2 Theoretical Background

To understand the extent to which agents might incorporate the incentive to gamble at a principal's expense, we implement the contest environment of Seel and Strack (2013) in the laboratory. In this section, we provide a brief review of the theoretical results from Seel and Strack (2013) that are relevant for this experimental implementation. Their basic setup is the following:

Each of two agents $i = 1, 2$ controls a project whose value is determined by a stochastic process $X^i = (X_t^i)_{t \in \mathbb{R}_+}$. The stochastic processes are governed by the following law of motion:

$$X_t^i = x_0 + \mu t + \sigma^2 B_t^i, \quad (1)$$

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_t^1 and B_t^2 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 is a stopping time, i.e., the decision when to irreversibly stop 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 players 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, that 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 arrives 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).

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

distribution

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

Thus, there is a unique prediction for the distribution of the stopped value.¹ Using the distribution F , Seel and Strack (2013) derive a closed-form solution for the expected value of each player's stopped process by computing $\int_0^\infty x dF(x)$. This enables them to prove the following main comparative statics result (taken from 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.*

We illustrate the two main propositions in Figures 1(a) and 1(b) using the parameters that we will use for our experimental implementation. Figure 1(b) shows the expected stopped value depending on the negative drift μ , while Figure 1(a) shows the induced cumulative distribution functions in equilibrium for the three drift parameters used in the experiment. The figures illustrate the main theoretical predictions: there is substantial amount of gambling (not stopping the process) in the contest although the process decreases in expectation.² Moreover, expected losses are non-monotone in the fundamentals.

Figure 1: The equilibrium distribution and expected stopped value for the parameters ($x_0 = 15, \sigma^2 = 2$) depending on the drift μ ; the implemented parameters are $\mu = -1.3, \mu = -0.24, \mu = -0.03$.

¹There are different, yet payoff-equivalent stopping strategies which lead to the equilibrium distribution F ; their characterisation is based on a result in probability theory by Skorokhod (1961).

²For comparison, without gambling, the cdf (not shown in the figure) is degenerate with all mass concentrated at $x_0 = 15$.

To get an intuition why contestants gamble 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, i.e., a profile in which both players stop immediately 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 $\frac{1}{2}$. Thus, the intuition of making the rival indifferent between her strategies is similar to simple games with mixed-strategy Nash equilibria such as Matching Pennies or Rock-Paper-Scissors. To implement the indifference condition in the contest game, the equilibrium distributions are strongly left-skewed to compensate for the downward drift of the process.

As the drift increases, the expected stopping time increases, but due to higher drift, the losses per unit of time decrease. In equilibrium, expected losses (expected stopping time times expected losses per unit of time) are largest if the drift is moderately negative.

Before we describe the experimental design in more detail, we state another property of the contest game which will be useful for the experimental implementation.

Remark 1 (*Independence of the Size of the Prize and the Risk Attitude*) Note that maximizing the expected payoff (prize times winning probability) yields the same result as maximizing the winning probability. Hence, 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. Thus, to obtain the main results, we only require that the expected utility of a player is increasing in her winning probability.

Remark 2 (*Skewed Stopped Value Distribution in Equilibrium*) The distribution function in Eq. (2) 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 experiment has two primary goals. The first is to test the prediction that contest incentives do indeed induce managers to delay shutting down projects that lose money in expectation. The second is to test the prediction that the expected losses from this gambling are not monotonic in the fundamentals. That is, in situations where project fundamentals are either close to even or very poor, the losses from gambling are small. However, for intermediate situations, these losses are larger.

The experimental design manipulates the expected change of the stochastic process across treatments. In addition, all treatments include a within-subject variation where the subjects are paid linearly in the value of their project. This variation provides a measure of every 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.

3.1 The contest game

The baseline implementation of the contest game uses the following setup: Two agents compete for a prize of 150 ECU.³ At every point in time, an agent privately observes the value of her own project. 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.⁴

For implementation purposes, we chose a discrete grid with size one and intervals of 0.25 seconds before the process changes its value. The implemented process is a random walk, where 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 $1 - p$. 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 A for the mapping from the discrete-time process parameters to the continuous-time process parameters. As the related literature (e.g., Anderson, Friedman, and Oprea, 2010, Pettit, Friedman, Kephart, and Oprea, 2014, and Oprea, 2014), we argue that the frequency of changes in the process is rapid enough

³Throughout, all payoffs are in experimental currency units, denoted ECU, which are converted at a rate of 0.01 Euros per ECU.

⁴Note that, in contrast to the benchmark theory, there is an exogenous upper bound on the length of the contest game. In the appendix, we provide several arguments why the time bound should not constrain predicted behaviour.

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. 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%.

3.2 Treatment design

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. The baseline value gives a moderately poor outlook for the project and, as will be detailed below, corresponds to the situation with the largest predicted losses in equilibrium. This value is the moderate-negative entry in Table 1. The extreme-negative value for the project fundamental gives an extremely poor outlook for the project, while the minimal-negative value means nearly zero expected losses. The extreme-negative and minimal-negative parameters were chosen to cover as broad a range of the parameter space as possible, while ensuring (at least to one-decimal place) the same theoretical prediction for the expected stopped value.⁵

Table 1: Summary of Between-Subject Treatment Conditions

Treatment	Abbreviation	Process Fundamental		Expected Stopped Value
		Drift (μ)	Pr(up)	
Minimal Negative	Min-	−0.03	0.49625	14.54
Moderate Negative (baseline)	Mod-	−0.24	0.47	13.71
Extreme Negative	Ext-	−1.3	0.3375	14.53

Notes: The drift parameter corresponds to the Brownian motion (continuous-time process); for all treatments the variance parameter is $\sigma = 2$ and the starting value is $X_0 = 15$. The Pr(up) value gives the probability of an increase in value for the associated random walk (discrete-time process); for all treatments the starting value is $X_0 = 15$, the time interval is $\Delta t = 0.25$ seconds and the jump size is $\Delta X = 1$.

⁵The drift parameter for the extreme-negative treatment could not be smaller without resulting in a theoretical prediction for the *discretised* version of the game that agents never gamble—that is, the discrete grid placed some restrictions on how low the drift could be, while maintaining in the support of the equilibrium strategy that players might wait for at least one uptick.

For every treatment, a within-subject variation is also included. This variation switches off the strategic element of the contest by paying a subject linearly in the stopped value of their project, reducing the environment to an individual choice setting. In this 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.*

3.3 Predictions

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

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 $\mu = -1.3$, 10.75 seconds for $\mu = -0.24$, and 30.66 seconds for $\mu = -0.03$.*

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 $\mu = -1.3$, 13.71 for $\mu = -0.24$, and 14.54 for $\mu = -0.03$.*

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

3.4 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.⁶

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 B 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).

⁶See Table 2 for a detailed summary of the treatments.

⁷Subjects 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.

Once everyone had correctly answered the comprehension questions, the session switched to ConG (Pettit, Friedman, Kephart, and Oprea, 2014), the main software interface in which the dynamic stopping task was implemented (see Appendix C for 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, 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 draw of the random walk described in Section 3.1 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 is 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. 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.

Table 2: Summary of treatments.

Treatment	Task Order	Number of		Earnings (Euro)		
		Matching Groups	Subjects	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

4 Main Experimental Results

This section provides the main experimental results, in particular the tests of the hypotheses stated in Section 3.3. Before these explicit tests, Figure 2 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. It is also clear that some form of gambling—that is stopping after time zero—takes place in all parts of the experiment.

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 in Figure 1(a), although the observed upper bound of the support is larger, particularly for the Min -ve and Mod -ve treatments.

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 (that is, did the subject gamble at all in a given period), the stopped time and the stopped value.⁸ The averages for percent gambling show that a significant proportion of subjects gamble for all drift parameters, in line with

⁸The main text focuses on the last five periods of the contest and lottery tasks. See Tables D.1 and D.2 for a comparison of these outcomes for the first five periods and over all periods.

(a) Stopped Time

(b) Stopped Value

Figure 2: The Observed Distribution of Stopped Times and Stopped Values

Table 3: Summary of the Outcomes in the Contest

Treatment	Percent Stopped Time > 0	Stopped Time (sec)		Stopped Value (ECU)	
		Average (95% C.I.)	Predicted	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. ***1%, **5%, *10% significance using cluster robust standard errors. Confidence intervals are constructed using cluster-robust standard errors.

Hypothesis 1A. The statistical significance for this is established using a random-effects probit regression on the probability of gambling 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.⁹

The stopped time column shows that the time spent gambling 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 9 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 gambling 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 D.3 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 average in the Mod -ve treatment being significantly smaller than in both the Min -ve and Ext -ve treatments, which have statistically similar

⁹This 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 level, 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 recent discussion of session-effects. Robustness checks using non-parametric tests on matching-group averages are reported in Tables 8 and 9 of the Appendix. They lead to the same conclusions.

averages—see Table 9 for further details and robustness checks. Indeed the U-shape prediction is more pronounced than predicted as 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 D.3 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: Summary of the Outcomes in the Lottery

Treatment	Percent Stopped Time > 0	Stopped Time (sec)		Stopped Value (ECU)	
		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. ***1%, **5%, *10% significance using cluster robust standard errors. Confidence intervals are constructed using cluster-robust standard errors. ***1%, **5%, *10% significance.

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 significant amount of gambling in all treatments, contrary to Hypothesis 3 (see also Table 8). Gambling 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 9 for these last two points). All point predictions for the lottery are rejected (see also Table D.3).

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.*

5 Further Analysis

To understand better the potential drivers of the primary treatment comparisons presented above, this section analyses subjects' behaviour in further detail. In the main text,

we focus on the extensive margin on gambling (percentage with stopping time greater than zero) with supplementary information on the intensity of gambling which we develop in more detail in the online appendix.

First, we look at how gambling develops with experience both in the contest and lottery tasks. We present evidence that experience and payoff structure play a differential role across treatments, in particular in the moderate negative and extreme negative conditions. Second, we present evidence for heterogeneous responses to the within-subject change in payoff structure. Finally, we pursue the idea of individual heterogeneity by investigating further the determinants of gambling choices under contest incentives.

5.1 Development of Behaviour over Periods

Figure 3(a) shows the observed pattern of gambling by subjects (that is, a stopped time strictly larger than zero) across periods, separated both by drift parameter (different panels) and by task order (different lines in a panel). The figure illustrates notable differences in the pattern of behaviour across the three drift parameter conditions, in particular with regard to how gambling develops over periods and responds to the payoff structure and task order. In the Min -ve treatment, very little changes over periods, and there is little difference in the propensity to gamble, which is on the whole very high, either between the contest or the lottery payoff conditions, or between the different order combinations. At the other end of the spectrum, there is a strong trend for less gambling with experience in the Ext -ve treatment. Here, for a given level of experience in terms of periods, subjects are less likely to gamble under the contest payoff structure than under the lottery payoff structure, especially for the contest-lottery order where there is a significant increase in the propensity to gamble following the switch to lottery payoffs. This effect is consistent with competition under contest payoffs driving out intrinsic motives for gambling.

The pattern of behaviour is different again in the Mod -ve treatment, which is the setting where subjects are confronted with the strongest incentives to gamble for strategic reasons under contest payoffs. Here, like the Min -ve treatment, there is no strong trend to reduce gambling over periods. Furthermore, there is no suggestion that gambling increases under the lottery payoffs as seen in the Ext -ve condition. Instead there is a reduction in gambling following the switch of regime in the contest-lottery order; under

(a) Observed

(b) Predicted

Figure 3: Behaviour across Periods by Task Order: Percentage with Stopped Time > 0

the lottery-contest order, gambling propensity remains roughly the same before and after the switch.

To quantify these observations, Table 5 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. Figure 3(b) shows the predicted behaviour from this regression model. Moving beyond the extensive margin, analogous linear random-effects regressions are used for (overall) stopped time and stopped value. These are reported in Table D.4 of the supplementary materials; Figures D.1 and D.2 show the average and predicted behaviour across periods.¹⁰

Table 5: Probability of Gambling across Periods Regressions.

	Min -ve		Mod -ve		Ext -ve	
Contest	0.02	(0.188)	0.08***	(0.004)	-0.26***	(0.001)
Pay Order	0.01	(0.623)	0.10**	(0.042)	-0.15**	(0.027)
Contest $\times$ Pay 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. Standard errors clustered at the matching group level. *** 1%, ** 5%, * 10%.

In the Ext-ve treatment there is a significant trend to reduce gambling with experience. For the Min-ve and Mod-ve treatments there is no such trend for the extensive margin, although overall time before stopping does reduce with experience in the Mod-ve. While experienced subjects in the Mod-ve treatment reduce their inclination to gamble in the lottery phase, neither their overall gambling time nor their stopped value is significantly different across parts. However, subjects that start with lottery incentives gamble in such a way that their average stopped value in this first part is lower than for the first part

¹⁰In 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 $\times$ 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 $\times$ pay order coefficients. Given subjects in part 2 had not experienced another incentive scheme, the first of theses gives the “purer” between-subjects comparison. For reference, Table D.5 in 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.

of the subjects who start with contest incentives. For the Ext-ve treatment, while there is no difference in the first-part average for stopped values, stopped values are lower for experienced subjects in the lottery phase compared to the contest phase, once the general trend for increasing stopped values has been controlled for.

5.2 Heterogeneity in Behaviour

Overall, there is evidence for heterogeneity in subjects' approach to the stopping problem, over and above the treatment variation in drift parameter and payment scheme. Few subjects never gambled throughout a part, except in the Ext-ve condition (where 19% and 14% never gambled in the contest and lottery parts, respectively). Table 6 reports subject-level averages calculated using just the last five rounds of either the contest or the lottery task, which are used to focus more on converged behaviour.¹¹ Perhaps unsurprisingly given the development of the overall averages seen in the previous section, many subjects gambled a majority of the time during the last five rounds of the Min-ve and Mod-ve conditions, while many gambled only a minority of the time in Ext-ve conditions. Nonetheless, there were non-negligible numbers of subjects doing something different for the latter two treatments—that is, gambling a minority of the time in the Mod-ve condition (16-17%) or gambling the majority of the time in the Ext-ve condition (21-34%). For the Min-ve condition, gambling the majority of the time was pervasive (over 95%).

Beyond the extensive margin, this subject heterogeneity is also seen in the way in which subjects that gambled stopped the process. While there were many that never used the stop now button, especially in the Ext-ve treatment where many never gambled at all, between 22% and 66% used it at least once. The threshold interfaces were more popular, and there is a more even split between those that used it the majority and those that used it a minority of times. By way of contrast, there is little heterogeneity in bankruptcy experiences, with the vast majority of subjects experiencing a bankruptcy in at most a minority of rounds. Furthermore, bankruptcy was less likely to be experienced

¹¹See Figures D.3-D.5 of the online appendix for the cumulative distribution of the subject-level averages (again using last five rounds data) for each of the following means of ending the stopping task: stopping during the warmup phase or immediate at the start of the active phase (that is, not gambling), stopping during the active phase by going bankrupt, stopping during the active phase using the stop now button, and stopping during the active phase by hitting one of the thresholds. The only outcome not included in the figures is never stopping, whereby the subject never stops the process during the active phase, but the value never hits bankruptcy; this happened rarely (see Table D.6). See Table D.7 for the subject-level averages calculated using all rounds of a part.

Table 6: Subject-Level Averages: Characterising Stopping Choices during the Last Five Rounds

Treatment	Contest				Lottery			
	Never 0 from 5	Minority 0-2 from 5	Majority 3-5 from 5	Always 5 from 5	Never 0 from 5	Minority 0-2 from 5	Majority 3-5 from 5	Always 5 from 5
<i>How often did subject not stop immediately (i.e. gamble time > 0)</i>								
Min-ve	0.00	0.04	0.96	0.88	0.01	0.04	0.96	0.85
Mod-ve	0.09	0.17	0.83	0.70	0.06	0.16	0.84	0.57
Ext-ve	0.45	0.79	0.21	0.05	0.42	0.66	0.34	0.17
<i>How often did subject stop using the stop button (during the active phase)</i>								
Min-ve	0.34	0.65	0.35	0.01	0.40	0.74	0.26	0.03
Mod-ve	0.42	0.74	0.26	0.01	0.46	0.85	0.15	0.01
Ext-ve	0.78	0.95	0.05	0.00	0.76	0.89	0.11	0.04
<i>How often did subject stop using threshold interface (during the active phase)</i>								
Min-ve	0.09	0.46	0.54	0.17	0.12	0.41	0.59	0.15
Mod-ve	0.19	0.53	0.47	0.20	0.16	0.46	0.54	0.16
Ext-ve	0.60	0.84	0.16	0.03	0.56	0.79	0.21	0.09
<i>How often did subject hit bankruptcy</i>								
Min-ve	0.70	0.99	0.01	0.00	0.57	0.95	0.05	0.00
Mod-ve	0.81	1.00	0.00	0.00	0.84	0.99	0.01	0.00
Ext-ve	0.99	1.00	0.00	0.00	0.97	1.00	0.00	0.00

Notes: The proportions of subjects are reported separately for the contest and lottery parts. Data from the last five rounds for each part.

in the more negative drift conditions, reinforcing the message that subjects' stopping strategy changes substantially in response to the underlying drift parameter.

To investigate heterogeneity in subjects' response to the within-subject payoff-structure variation, not just their general inclination towards gambling, we expand the basic random-effects probit model for gambling, used in the previous subsection, to add random coefficients for the contest term and its interaction with pay order. The analysis of trend and incentive-scheme effects is robust to the inclusion of random coefficients—see Table D.8 of the supplementary materials for details. However, a likelihood ratio test of the random coefficient terms shows that they are strongly significant ($p < 0.001$) for all drift-parameter treatments, compared to the standard model where these coefficients are fixed. This result provides evidence of a significant degree of heterogeneity in subjects' response to the competitive incentives of the contest, and suggests a degree of heterogeneity in their strategic approach or reasoning.

5.3 Determinants of Gambling in the Contest

Stepping back from heterogeneity in the response to the within-subject variation, here we investigate further the heterogeneity in subjects' general inclination towards gambling under the contest incentive scheme. This analysis uses the observed behaviour from the lottery periods to construct an individual measure of a subject's preference for gambling when faced with the same stochastic process as in the contest periods. This individual measure and an indicator for whether the subject gambled in the first match of the contest are used to explicitly model subject gambling types. Model 1 of Table 7 shows the results of a regression analysis of the probability of gambling in the contest matches using these individual measures and controls for the task order and the match number. In all treatments these individual components play a significant role in the probability of gambling in a contest period. For the Min-ve and Mod-ve, it is the lottery gambling average that plays the most important role; for the Ext-ve treatment, it is the indicator for gambling in the first contest match.

Model 2 of Table 7 adds as explanatory variables whether a subject's opponent from the previous match had a stopped value equal to 15 or strictly greater than 15. That is, the second model investigates the role played by the revealed gambling choices of opponents from the previous period. For all treatments, observing their previous match to have gambled increases the chance that a subject gambles, while observing them with a stopped value equal to the start value reduces the chance that a subject gambles. However it is only in the Mod-ve condition that this feedback has a significant bearing on subjects' choice to gamble.

Overall, the results of this analysis complement the patterns of behaviour illustrated in Figure 3: In the Min-ve treatment—where the cost of gambling is not high and the feedback from the stochastic process noisy (i.e. it is very close to a random walk without drift)—gambling rates are high, there is less of a role for individual heterogeneity and there is little evidence of a learning trend either over matches, order variation or in response to the other player's outcome, at least in the extensive margin (that is, the decision to not stop immediately). For the Ext-ve treatment—where the cost of gambling is high and the feedback from the stochastic process clear—there is a strong trend over matches to reduce gambling, and the initial choice to gamble in the first contest match is the most important individual component. It is in the intermediate Mod-ve treatment—

where subjects are most exposed to the inefficient incentives to gamble introduced by the contest payment scheme—that subjects respond significantly to their opponent’s outcome from the previous match.

Table 7: Determinants of Gambling Regressions.

	Min -ve		Mod -ve		Ext -ve	
<i>Model 1</i>						
Lottery-Contest Order	−0.00	(0.864)	−0.06**	(0.046)	−0.16**	(0.023)
Match	0.00	(0.275)	−0.00	(0.504)	−0.04***	(0.000)
Gambles in Match 1	0.07*	(0.081)	0.03	(0.561)	0.30***	(0.000)
Lottery Gambling Avg.	0.09**	(0.034)	0.33***	(0.000)	0.11	(0.221)
<i>Model 2</i>						
Lottery-Contest Order	0.00	(0.999)	−0.06**	(0.039)	−0.15**	(0.030)
Match	0.00	(0.269)	−0.00	(0.632)	−0.03***	(0.000)
Gambles in Match 1	0.08*	(0.067)	0.03	(0.523)	0.30***	(0.000)
Lottery Gambling Avg.	0.10**	(0.027)	0.31***	(0.000)	0.11	(0.212)
Other's Last Stopped Value = 15	−0.02	(0.136)	−0.04**	(0.019)	−0.03	(0.344)
Other's Last Stopped Value > 15	0.01	(0.535)	0.04**	(0.022)	0.05	(0.512)

Notes: Correlated random-effects probit models. The second model is the same as the first model except that it includes indicators for the other player’s outcome from the previous period. Coefficients report the marginal effect on the dependent variable. Standard errors clustered at the matching group level. *** 1%, ** 5%, * 10%.

The following results collect the main findings from this section.

Result 4 *In the Ext-ve treatment there is a significant trend to reduce gambling with experience. For the Min -ve and Mod -ve treatments there is no such trend for the extensive margin. In the Ext-ve treatment, contest payoffs appear to crowd out intrinsic motives to gamble. The same is not true in the Mod-ve treatment, with evidence that subjects are less inclined to gamble under lottery incentives.*

Result 5 *In all treatments, there is evidence for individual heterogeneity in the propensity to gamble. There is also evidence for heterogeneity in the response to the change in payoff structure.*

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

6 Alternative Explanations

This section considers two alternative explanations for the observed behaviour in our experiments. Specifically, whether adding concerns for regret or a joy-of-gambling to subjects’ decision-making, over and above maximising the expected utility of winning the prize, might explain our results. Regret is a natural candidate for an alternative explanation for two reasons. First it has been studied in the theoretical literature for related decision environments (Feng and Hobson, 2016b; Strack and Viefers, 2019). Second, the software interface suggests a potential role for regret: subjects see the complete realisation of the stochastic process, even if they “stopped” the process earlier in a period.¹² Including an intrinsic motivation for gambling also has an established theoretical literature (Diecidue, Schmidt, and Wakker, 2004); it is also a natural analog to including a joy-of-winning, which has been reported in the experimental literature on contests, in particular all-pay auctions and Tullock contests (see, for example, the survey by Dechenaux, Kovenock, and Sheremeta, 2015).¹³

These alternatives are judged on their ability to explain gambling behaviour under both lottery and contest incentives, as well as the observation that contest incentives crowd out gambling in the Ext-ve treatment, while the opposite is true in the Min-ve treatment. Explaining this latter observation will require considering subjects’ reasoning about their match’s strategy under contest payoffs.

6.1 Regret

A subject might regret to have stopped if he does not win the prize and—in hindsight—there was a winning strategy. In our implementation, subjects observe the whole sample path of the process, i.e., a feeling of regret due to a failure to stop at the “right” time seems quite natural. If such feeling of regret is not anticipated by the subject, behavioural predictions do not change compared to the main text. However, if the subject anticipates his regret, behavioural predictions might change. This type of anticipated regret was first

¹²Note that, this design choice was made to ensure that subjects gain the same amount of experience with the Brownian motion irrespective of their stopping strategy.

¹³Note that simply adding an intrinsic motive for winning would not change any predictions in risk contest implemented in our experiments—as noted in Remark 1, predicted behaviour is independent of the utility agents gain from winning the prize, so long as expected utility increases linearly in the probability of winning.

studied in Loomes and Sugden (1982) and recently extended to dynamic settings.

In particular, the literature considers two natural notions of regret for optimal stopping problems: regret over stopping below the past peak (i) until the stopping time or (ii) until the deadline T . Strack and Viefers (2019) characterize optimal behaviour for the two types of regret in a related individual choice setting and also implement their choice problem in the laboratory. Their setting differs in three aspects from our individual choice setting. First, they implement a geometric Brownian motion rather an arithmetic Brownian motion and thus have no risk of bankruptcy. Second, the payoff is given by the stopped value rather than influencing the probability of winning a fixed prize. Finally, they include a breakdown probability; that is, after each unit of time, there is a positive probability that the decision problem ends with a zero payoff for the subject.

Regret over stopping below the peak before the deadline is shown to be empirically indistinguishable from the standard theory without regret in Strack and Viefers (2019). Their arguments extend to our setting. Thus, regret over stopping below the peak until the deadline offers no help in explaining the gap between the theory and the data for the individual choice treatment.

Strack and Viefers (2019) report evidence from their experiment in favour of the most popular theory of regret, regret over stopping below the past peak until the stopping time. For instance, this theory offers a possible explanation for the disposition effect in behavioural finance (Shefrin and Statman, 1985)—i.e., traders sell winning assets while keeping losing assets in their portfolio. In our individual choice setting, an individual who feels this type of regret should stop immediately: this behaviour is optimal without regret and makes sure that no regret is felt. Thus, Proposition 3 directly extends. Intuitively, there is now even more pressure to stop earlier. Clearly, this formalisation of regret is also not able to narrow the gap between the theory and the data in our individual choice setting.¹⁴

¹⁴In our individual choice setting, agents receive information about whether they won or not, but not about the stopping value which would have been sufficient to win. There would be different ways to adapt the regret notion to this setting. For all of these adaptations, however, following the argumentation in the text, an explanation based on regret over stopping below the past maximum can be rejected.

6.2 Joy of Gambling

There is empirical evidence that people might accept or reject gambles because they derive an extra procedural (dis-)utility, which depends on the circumstance and properties of the gamble (Le Menestrel, 2001). For instance, otherwise risk-averse individuals sometimes gamble on roulette spins, horse races and other sports. Similarly, subjects in our experiment might enjoy watching the process with some stake on the line and therefore delay their stopping decision. Diecidue, Schmidt, and Wakker (2004) show that this procedural utility of gambling is inconsistent with the standard framework of expected utility theory. Instead, the related literature separates two terms, a term for the gamble which is evaluated according to expected utility theory and an additional term which measures the intrinsic cost or benefit of not gambling. We verbally discuss two models below and provide the corresponding derivations in Appendix A.4.

The simplest way to include joy of gambling in our model is by an additional constant gain per unit of time that a subject gambles. Under this specification, the predictions for the individual choice treatment are: an individual either chooses not to gamble at all or he gambles until he is bankrupt or the value of the process reaches at least 30. Thus, a specification with constant joy of gambling is not able to explain observed behaviour.

Instead, suppose that the utility of gambling is decreasing over time. Intuitively, it might be fun to play the stopping task for a couple of seconds, but it might become tiring after a while. Such a model allows for richer patterns in the data. In the individual choice treatment, the expected decrease in probability of winning is constant per unit of time is constant. Hence, if the extra utility of gambling decreases over time, an individual might have a positive, but short optimal stopping time. Intuitively, contestants value a thrill in the beginning, but the effect of a decreasing winning probability becomes dominant thereafter.

6.3 Reasoning About Opponents

In the Ext-ve treatment, subjects are more inclined to stop immediately under contest incentives compared to lottery incentives. To accommodate this pattern of behaviour, we suggest the following pattern of reasoning: subjects observe that most rivals stop immediately. Thus, by stopping immediately, the probability of winning the contest is 50%. While they are not able to determine the best response to an opponent who stops

immediately—which is to gamble, stopping the first time the process reaches 16 or 0—they realize that continuing for an additional unit of time at the initial value results in a huge decrease in winning probability. In particular, while they only risk a decrease in winning probability by $\frac{1}{30}$ in the individual choice treatment, the same decision reduces the winning probability by $\frac{1}{2}$ in the contest when the rival stops at 15 and the own process decreases to reach 14. Consequently, they gamble less in the contest than the lottery, and competition crowds out the intrinsic joy of gambling.

In the moderate negative treatment, however, we observe more gambling in the contest. Intuitively, individuals do not observe that their rivals stop at 15 as often. Thus, they do not have a point belief about their opponent’s stopped value anymore—supported by the observation that subjects’ gambling decision is significantly affected by the revealed stopped value of the opponent from the previous contest period. In turn, this makes gambling more attractive at the starting value. Finally, in the minimal negative treatment, the loss in expected value is so small that individuals gamble quite a bit both in the individual choice and in the contest treatment. To sum up, the observed data is consistent with decreasing joy of gambling combined with myopic reasoning whether to stop.

6.4 Skewing the Stopped Value Distribution

One feature of the predicted equilibrium behaviour for the contest is that players gambling for strategic purposes should employ a win-small-lose-big strategy. That is, under contest incentives, they should be willing to risk large losses for small gains. This type of stopping strategy distorts the stopped value distribution so that it is left-skewed. In contrast, intrinsic motives for gambling that are tied to the amount of time spent actively gambling would not have this feature (see Remark 2, and Remark 3 in the online appendix for the decreasing joy of gambling model). Consequently, the skewness of the stopped value distribution can be used to further distinguish between gambling motives.

This analysis is of particular interest for the Mod-ve treatment. It is this treatment that most exposes subjects to the inefficient incentives to gamble for strategic rather than intrinsic motives. For example, the between-sessions contrast for this treatment—which compares the behaviour of subjects under contest incentives versus those under lottery incentives—found that, while the inclination towards gambling and the amount of time spend gambling was on average very similar in part 2, subjects seemed to gamble in a

different manner so that the stopped value was on average lower under lottery incentives than under contest incentives. This difference in gambling strategy can be seen in the skewness statistics (see Table D.10 of the online appendix for details): In the standard order, the stopped value distribution is negatively skewed under the contest incentives (marginally significantly different from zero for all rounds of part 2, and significantly different for the last five rounds). However, for the reverse order, the skewness of the stopped value distribution is not significantly different from zero under lottery incentives (when looking at either all rounds or the last five rounds of part 2). Furthermore, for these reverse order sessions, the stopped value distribution is subsequently significantly negatively skewed under the contest incentives of part 3. This is most cleanly illustrated by comparing behaviour from period 15 (last round of lottery incentives) with period 16 (first round of contest incentives), where the former is negative but not statistically different from zero, but the latter is more negative and statistically different from zero.

Finally, the distribution of stopped values (see Figure D.7 of the online appendix) suggests a further difference in stopping strategy between the lottery and contest decisions of the Min-ve treatment. In the lottery, there is some evidence that subjects are relatively more inclined to stop at a round number (such as 5 or 10), something that does not seem a relevant consideration under contest incentives (p-value using a regression-based test is 0.037 for all rounds and 0.012 for the last five rounds).

7 Conclusion

This paper presents a theory-based experimental investigation of behaviour in risk-taking contests using a novel laboratory stopping task. Despite its continuous-time 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. It maintains a tight connection to the underlying theory. Further, predicted behaviour is independent of the risk attitude of the participants and every single round of the experiment can be compensated without distorting incentives. This implementation of choices over probability distributions through a privately observed stochastic process 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—excessive risk-taking and expected losses which are non-monotone in the drift—are observed in the experiment. We find evidence for excessive 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 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 pay subjects linearly in their 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 excessive risk-taking appear to reinforce the non-monotonicity in subtly different ways: Under the extreme negative condition, contest incentives appear to crowd out the intrinsic motive for excessive-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 excessive 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 to gamble—but it is only a starting point towards improving behavioural predictions for risky environments by combining empirical evidence with theoretical models.

References

- ANDERSON, S. T., D. FRIEDMAN, AND R. OPREA (2010): “Preemption Games: Theory and Experiments,” *American Economic Review*, 100, 1778–1803.
- ANKIRCHNER, S., AND P. STRACK (2011): “Skorokhod Embeddings In Bounded Time,” *Stochastics and Dynamics*, 11, 215–226.
- BERNINGHAUS, S. K., K.-M. EHRHART, AND M. OTT (2006): “A Network Experiment

- in Continuous Time: The Influence of Link Costs,” *Experimental Economics*, 9, 237–251.
- (2012): “Forward-Looking Behavior in Hawk-Dove Game in Endogenous Networks: Experimental Evidence,” *Games and Economic Behavior*, 75, 35–52.
- BIGONI, M., M. CASARI, A. SKRZYPACZ, AND G. SPAGNOLO (2015): “Time Horizon and Cooperation in Continuous Time,” *Econometrica*, 83, 587–616.
- BISMUT, J.-M., AND B. SKALLI (1977): “Temps d’arrêt optimal, théorie générale des processus et processus de Markov,” *Probability Theory and Related Fields*, 39, 301–313.
- CASON, T. N., D. FRIEDMAN, AND E. HOPKINS (2014): “Cycles and Instability in a Rock-Paper-Scissors Population Game: A Continuous Time Experiment,” *Review of Economic Studies*, 81, 112–136.
- CHARNESS, G., AND P. KUHN (2011): “Lab labor: What can labor economists learn from the lab?,” in *Handbook of labor economics*, vol. 4, pp. 229–330. Elsevier.
- CHEVALIER, J., AND G. ELLISON (1997): “Risk Taking by Mutual Funds as a Response to Incentives,” *Journal of Political Economy*, 105, 1167–1200.
- (1999): “Carrier Concerns of Mutual Fund Managers,” *The Quarterly Journal of Economics*, 114, 389–432.
- DECHENAUX, E., D. KOVENOCK, AND R. M. SHEREMETA (2015): “A survey of experimental research on contests, all-pay auctions and tournaments,” *Experimental Economics*, 18, 609–669.
- DECK, C., AND N. NIKIFORAKIS (2012): “Perfect and imperfect real-time monitoring in a minimum-effort game,” *Experimental Economics*, 15, 71–88.
- DIECIDUE, E., U. SCHMIDT, AND P. P. WAKKER (2004): “The Utility of Gambling Reconsidered,” *Journal of Risk and Uncertainty*, 29, 241–259.
- FALKENSTEIN, E. G. (1996): “Preference for Stock Characteristics as Revealed by Mutual Fund Portfolio Holdings,” *Journal of Finance*, 51, 111–135.

- FANG, D., AND T. NOE (2016): “Skewing the Odds: Taking risks for rank-based rewards,” *Working Paper*.
- (2018): “Less competition, more meritocracy?,” *Working Paper*.
- FANG, D., T. NOE, AND P. STRACK (2020): “Turning up the heat: The discouraging effect of competition in contests,” *Journal of Political Economy*, 128(5), 1940–1975.
- FENG, H., AND D. HOBSON (2015): “Gambling in contests modelled with diffusions,” *Decisions in Economics and Finance*, 38(1), 21–37.
- (2016a): “Gambling in Contests with Random Initial Law,” *Annals of Applied Probability*, 26, 186–215.
- (2016b): “Gambling in Contests with Regret,” *Mathematical Finance*, 26, 674–695.
- FISCHBACHER, U. (2007): “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 10(2), 171–178.
- FRÉCHETTE, G. (2011): “Session Effects in the Laboratory,” *Experimental Economics*, 15, 485–498.
- FRIEDMAN, D., AND R. OPREA (2012): “A continuous dilemma,” *American Economic Review*, 102, 337–363.
- GREINER, B. (2004): “The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics,” *Working Paper*.
- (2015): “Subject pool recruitment procedures: organizing experiments with ORSEE,” *Journal of the Economic Science Association*, 1(1), 114–125.
- HART, S. (2008): “Discrete Colonel Blotto and General Lotto Games,” *International Journal of Game Theory*, 36, 441–460.
- HUANG, J., C. SIALM, AND H. ZHANG (2011): “Risk Shifting and Mutual Fund Performance,” *Review of Financial Studies*, 24, 2575–2616.
- HVIDE, H. (2002): “Tournament Rewards and Risk Taking,” *Journal of Labor Economics*, 20, 877–898.

- KEMPF, A., AND S. RUENZI (2008): “Tournaments in Mutual-Fund Families,” *Review of Financial Studies*, 21, 1013–1036.
- KEPHART, C., AND D. FRIEDMAN (2015): “Hotelling revisits the lab: equilibration in continuous and discrete time,” *Journal of the European Science Association*, 1, 132–145.
- KIRCHLER, M., F. LINDNER, AND U. WEITZEL (2018): “Rankings and Risk-Taking in the Finance Industry,” *Journal of Finance*, 73, 2271–2302.
- LAZEAR, E. P., AND S. ROSEN (1981): “Rank-Order Tournaments as Optimal Labor Contracts,” *Journal of Political Economy*, 89, 841–864.
- LE MENESTREL, M. (2001): “A Process Approach to the Theory of Gambling,” *Theory and Decision*, 50, 249–262.
- LIZZERI, A., AND N. PERSICO (2001): “The Provision of Public Good under Alternative Electoral Incentives,” *American Economic Review*, 91, 225–239.
- LOOMES, G., AND R. SUGDEN (1982): “Regret Theory: An Alternative Theory of Rational Choice under Uncertainty,” *Economic Journal*, 92, 805–824.
- MAGNANI, J., A. GORRY, AND R. OPREA (2016): “Time and State Dependence in an Ss Decision Experiment,” *American Economic Journal: Macroeconomics*, 8, 283–310.
- MYERSON, R. B. (1993): “Incentives to Cultivate Favored Minorities under Alternative Electoral Systems,” *American Political Science Review*, 87, 856–869.
- NIEKEN, P., AND D. SLIWKA (2010): “Risk-taking tournaments - Theory and Experimental Evidence,” *Journal of Economic Psychology*, 31, 254–268.
- OPREA, R. (2014): “Survival versus Profit Maximization in a Dynamic Stochastic Experiment,” *Econometrica*, 82, 2225–2255.
- OPREA, R., D. FRIEDMAN, AND S. T. ANDERSON (2009): “Learning to Wait: A Laboratory Experiment,” *Review of Economic Studies*, 76, 1103–1124.
- OPREA, R., K. HENWOOD, AND D. FRIEDMAN (2011): “Separating the Hawks from the Doves: Evidence from continuous time laboratory games,” *Journal of Economic Theory*, 146, 2206–2225.

- PESKIR, G., AND A. SHIRYAEV (2006): *Optimal Stopping and Free-Boundary Problems*. Birkhäuser Verlag, Basel, Boston, Berlin.
- PETTIT, J., D. FRIEDMAN, C. KEPHART, AND R. OPREA (2014): “Software for continuous game experiments,” *Experimental Economics*, 17, 631–648.
- SEEL, C. (2015): “Gambling in Contests with Heterogeneous Loss Constraints,” *Economics Letters*, 136, 154–157.
- SEEL, C., AND P. STRACK (2013): “Gambling in Contests,” *Journal of Economic Theory*, 148, 2033–2048.
- (2016): “Continuous time contests with private information,” *Mathematics of Operations Research*, 41(3), 1093–1107.
- SHEFRIN, H., AND M. STATMAN (1985): “The Disposition to Sell Winners too Early and Ride Losers Too Long: Theory and Evidence,” *Journal of Finance*, 40, 777–790.
- SKOROKHOD, A. V. (1961): “Issledovaniya po teorii sluchainykh protsessov (Stokhasticheskie differentsialnye uravneniya i predelnye teoremy dlya protsessov Markova),” *Izdat Kiev University*.
- STRACK, P. (2016): “Risk-taking in Contests - The Impact of Fund-Manager Compensation on Investor Welfare,” *Working Paper*.
- STRACK, P., AND P. VIEFERS (2019): “Too Proud to Stop: Regret in Dynamic Decisions,” *Journal of the European Economic Association*.

Appendix

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 $\min \left\{ \frac{X_t}{30}, 1 \right\} = \frac{1}{30} \min \{X_t, 30\}$. Thus, the player chooses a stopping time τ to solve

$$\max_{\tau} \frac{1}{30} \min \{X_{\tau}, 30\}.$$

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

$$\mathbb{E}\left(\frac{1}{30} \min \{X_\tau, 30\}\right) \leq \frac{1}{30} \mathbb{E}(X_\tau) \leq \frac{1}{30} X_0 = \frac{1}{2}.$$

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.

Robustness Checks for Main Hypothesis Tests

Table 8: Probability Stopped Time > 0 : Robustness Checks using Non-Parametric Tests

Treatment	Percent Stopped	Regression based	K-S p-value based on:		
	Time > 0	p-value	All data	Subject	Matching group
<i>Hypothesis 1a: Probability stopped time > 0 in the contest</i>					
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
<i>Hypothesis 3: Probability stopped time > 0 in the lottery</i>					
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 either 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 put all weight on stopping immediately and the observed distribution using either all the data, subject averages or matching group averages. All tests use data from the last five rounds of the contest and the lottery parts, respectively.

Table 9: Treatment Comparisons in Hypotheses 1–3: Robustness Checks using Non-Parametric Tests

Comparison	Regression based	Rank sum	K-S p-value based on:		
	p-value	p-value	All data	Subject	Matching group
<i>Hypothesis 1b: Stopped time in the contest</i>					
Min -ve vs Mod -ve	0.000	0.000	0.000	0.000	0.001
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
<i>Hypothesis 2a: Stopped value in the contest</i>					
Min -ve vs Mod -ve	0.003	0.008	0.000	0.052	0.055
Min -ve vs Ext -ve	0.989	0.289	0.000	0.418	0.400
Mod -ve vs Ext -ve	0.000	0.000	0.000	0.000	0.001
<i>Hypothesis 3: Stopped time in the lottery</i>					
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
<i>Hypothesis 3: Stopped value in the lottery</i>					
Min -ve vs Mod -ve	0.002	0.010	0.000	0.052	0.055
Min -ve vs Ext -ve	0.759	0.623	0.000	0.418	0.400
Mod -ve vs Ext -ve	0.000	0.002	0.000	0.002	0.003

Notes: The regression-based p -values are based on a linear random-effects regression (with clustering at the matching group level) of either 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 rounds of the contest and the lottery parts, respectively.

Contents

A Further Details on the Contest Game	3
A.1 Discretisation	3
A.2 Point Predictions	3
A.3 Analytical Time Bound	4
A.4 Introducing Joy of Gambling	5
B Example Instructions	8
C Screenshots	14
D Further Results	16
D.1 Additional Tables	16
D.2 Additional Figures	21

List of Figures

C.1 Screenshot: Warm-up Phase	14
C.2 Screenshot: Active Phase	14
C.3 Screenshot: Feedback Phase	15
D.1 Behaviour across Periods by Task Order: Stopped Time	21
D.2 Behaviour across Periods by Task Order: Stopped Value	22
D.3 Proportion of Rounds Subject Gambled (Stopped Value > 0) during the Last Five rounds.	23
D.4 Proportion of Rounds Subject Gambled and Stopped using the Button Interface during the Last Five rounds.	23
D.5 Proportion of Rounds Subject Gambled and Stopped using Threshold In- terface during the Last Five rounds.	24
D.6 Proportion of Rounds Subject Gambled and went Bankrupt during the Last Five rounds.	24
D.7 Stopped Value Distribution in the Mod-ve Treatment	25

List of Tables

D.1	Summary of the Outcomes in the Contest: Comparison across Data Sub-samples	16
D.2	Summary of the Outcomes in the Lottery: Comparison across Data Sub-samples	16
D.3	p -Values from Point-Prediction Tests for Hypotheses 1–3	17
D.4	Random-Effects Regressions of Behaviour across Periods: Probability of Gambling, Stopped Time and Stopped Value.	17
D.5	Complete Set of Contrast Tests from Regressions on Behaviour across Periods	18
D.6	Aggregate-Level Stopping Choices across All Rounds	18
D.7	Subject-Level Averages: Characterising Stopping Choices across All Rounds	19
D.8	Heterogeneity in Response to change in Payoff Structure: Probability of Gambling across Periods	19
D.9	Further Statistics on the Distribution of Stopped Values	20
D.10	Stopped Value Variance and Skewness in the Mod-ve Treatment	20

A Further Details on the Contest Game

A.1 Discretisation

The approach in the discretisation in this paper is similar to that of Anderson, Friedman, and Oprea (2010). However, there is a minor difference, because we consider an arithmetic Brownian motion rather than a geometric Brownian motion. By the functional central limit theorem, the random walk considered in the implementation converges in distribution to a Gaussian process (here: the Brownian motion with drift μ and variance σ^2). The details are given below:

Denote by Δx the jump size of the random walk, by Δt the time between the jumps, and by p the uptick probability. Assume $n = \frac{t}{\Delta t}$ is an integer. Thus, the expected value of the random walk R at time t is $\mathbb{E}[R(t)] = n(\Delta x)(2p - 1)$ and the variance is $\text{var}[R(t)] = n(\Delta x)^2(1 - (2p - 1)^2)$.

Let $\Delta x = \sigma\sqrt{\Delta t}$ and $p = [1 + \frac{\mu}{\sigma}\sqrt{\Delta t}]/2$. Then, for $\Delta \rightarrow 0$, $\mathbb{E}[R(t)] \rightarrow \mu t$ and $\text{var}[R(t)] \rightarrow \sigma^2 t$. Thus, in the limit, expected value and variance converge to the parameters in the Brownian motion setting.

In our implementation, we have $\Delta t = 0.25$ seconds and $\sigma = 2$. Thus, we obtain $\Delta x = 1$ and $p = 0.3375$ for $\mu = -1.3$, $p = 0.47$ for $\mu = -0.24$ and $p = 0.49625$ for $\mu = -0.03$, which are the values we have used for the different treatments.

A.2 Point Predictions

Note that according to Seel and Strack (2013), p.2043, the expected value of the stopped processes in equilibrium satisfies

$$\mathbb{E}(X_\tau) = \frac{\sigma^2}{2\mu} + \left(1 + \frac{1}{2(\exp(-\frac{2\mu x_0}{\sigma^2}) - 1)}\right) \left(x_0 - \frac{\sigma^2 \log(2 - \exp(\frac{2\mu x_0}{\sigma^2}))}{2\mu}\right).$$

Plugging in $\sigma = 2$, $x_0 = 15$ and the different values of μ , we obtain the values in Hypothesis 2(b).

If we divide the expected losses $(15 - \mathbb{E}(X_\tau))$ by the losses per period, we obtain the average amount of periods until stopping occurs. Since a period takes 0.25 seconds, we

can multiply this result by 4 to obtain the expected stopping time in seconds, i.e., the point prediction in Hypothesis 1(c).

A.3 Analytical Time Bound

In this section, we provide a calculation of an analytical time bound such that the equilibrium distribution can be implemented with a strategy that stops up to that time for the most critical treatment (min-ve). We used this time bound to determine the length of each round of the experiment.

We calculate the time bound T such that the two-player equilibrium distribution derived in Seel and Strack (2013) can be implemented before time T . For this purpose, recall the distribution in the paper is given by

$$F(x) = \min \left\{ \frac{1}{2} \frac{\exp(\frac{-2\mu x}{\sigma^2}) - 1}{\exp(\frac{-2\mu x_0}{\sigma^2}) - 1}, 1 \right\} \quad (3)$$

for $x \geq 0$.

For the condition which we are going to use we need the probability distribution for a centered process, i.e., a process which starts at zero. We obtain this by shifting the distribution to the left:

$$F(x) = \min \left\{ \frac{1}{2} \frac{\exp(\frac{-2\mu(x+x_0)}{\sigma^2}) - 1}{\exp(\frac{-2\mu x_0}{\sigma^2}) - 1}, 1 \right\} \quad (4)$$

for $x \geq -x_0$.

We are now ready to state the condition for a bounded time embedding of a distribution given in Ankirchner and Strack (2011). They define $g(x) = F^{-1}(\Phi(x))$, where $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp(\frac{z^2}{2}) dz$ is the density function of the normal distribution.

Lemma 1 (Ankirchner and Strack, 2011, Theorem 2) *Suppose that $g(\cdot)$ is Lipschitz-continuous with Lipschitz constant $\sqrt{T}$. Then the distribution F can be embedded in $X_t = \mu t + B_t$, with a stopping time that stops almost surely before T .*

Hence, we need the inverse of the equilibrium distribution in Eq. (4) on its support to verify the condition in Ankirchner and Strack (2011):

$$F^{-1}(x) = -\frac{\sigma^2}{2\mu} \ln(1 + 2x(\exp(\frac{-2\mu x_0}{\sigma^2}) - 1)) - x_0 \quad (5)$$

We take the derivative to get

$$(F^{-1})'(x) = -\frac{\sigma^2}{2\mu} \frac{2(\exp(\frac{-2\mu x_0}{\sigma^2}) - 1)}{1 + 2x(\exp(\frac{-2\mu x_0}{\sigma^2}) - 1)} \quad (6)$$

We show Lipschitz continuity of $F^{-1}(x)$ and derive the corresponding Lipschitz constants.

Since $\mu < 0$, $(F^{-1})'(x)$ is positive and decreasing in x for $x \in [0, 1]$. Thus,

$$(F^{-1})'(x) \leq (F^{-1})'(0) = -\frac{\sigma^2}{\mu} (\exp(\frac{-2\mu x_0}{\sigma^2}) - 1) \quad (7)$$

for all $x \in [0, 1]$. Thus, $(F^{-1})'(x)$ is Lipschitz continuous with Lipschitz constant $-\frac{\sigma^2}{\mu} (\exp(\frac{-2\mu x_0}{\sigma^2}) - 1)$. The normal distribution is Lipschitz continuous with Lipschitz constant $\frac{1}{\sqrt{2\pi}}$, where π is the mathematical constant. Thus, $g(x)$ is Lipschitz continuous with Lipschitz constant $-\frac{\sigma^2}{\mu} (\exp(\frac{-2\mu x_0}{\sigma^2}) - 1) \frac{1}{\sqrt{2\pi}}$. From Lemma 1, we obtain that F can be embedded if

$$T \geq (-\frac{\sigma^2}{\mu} (\exp(\frac{-2\mu x_0}{\sigma^2}) - 1) \frac{1}{\sqrt{2\pi}})^2. \quad (8)$$

For the min-ve treatment ($\mu = -0.03, \sigma = 2, x_0 = 15$), the time bound is roughly 90 seconds, the duration of each round of the experiment. Intuitively, a constraint through the duration is more likely to change behaviour for small values of μ , where there is more incentive for gambling. Thus, our choice of deadline should not influence behaviour since it respects the bounded time condition for the most critical treatment.

A.4 Introducing Joy of Gambling

Diecidue, Schmidt, and Wakker (2004) propose a general model which allows to capture joy of gambling, but is also able to accomodate an intrinsic disutility of a static gamble. They call it the holistic cost of the gamble. Adapted to our dynamic setting, we can study such a holistic cost over time with an additional term $C(X_\tau) = \int_0^\tau c(t, X_t) dt$. We now propose two models which impose more assumption on the cost term.

Model 1: Constant Joy of Gambling

In the most simple model, the cost of absence of gambling would be constant over time and independent of the value of the process, i.e., $C(X_\tau) = \int_0^\tau c dt = c\tau$. Thus, in the

individual choice treatment, a subject with constant joy of gambling aims to maximize

$$\frac{X_\tau}{30} + c\tau.$$

The following lemma states the behavioural predictions for such a subject.

Lemma 2 *There is a threshold $\bar{c}$ such that for all $c < \bar{c}$, stopping immediately is optimal and for all $c > \bar{c}$, the subject should not stop in $(0, 30)$.*

Since c is constant, the expected value of continuing for an additional unit of time is the same for all $X_t \in (0, 30)$. The expected value is increasing in c . Let $\bar{c}$ be such that the expected value of continuation is negative for $c < \bar{c}$ and positive for $c > \bar{c}$. Then the process is a supermartingale (submartingale) for $c < \bar{c}$ ($c > \bar{c}$) and the proof of Proposition 3 directly extends to the result in Lemma 2.

Thus, for generic parameters of c , the linear model predicts that stopping in the individual choice treatment only occurs at 0, 15, and at or above 30.

Model 2: Decreasing Joy of Gambling

Now suppose the holistic cost of absence of gambling is decreasing over time, $C(X_\tau) = \int_0^\tau c(t)dt$, where $c(t)$ is decreasing in t and, for simplicity, it does not depend on X_t . The expected loss in winning probability is constant over time. Thus, the gain of continuation decreases, i.e., there is an optimal threshold for the stopping time τ . Moreover, note that (i) the threshold value increases as the drift increases and (ii) the optimal stopping decision can be based on whether it is profitable to continue for one additional unit of time only. Prediction (i) is consistent with the data (see Result 3) and prediction (ii) says that a myopic decision-making leads to an overall optimal decision.

When $c(t)$ does not depend on X_t and the process is stopped at a given time t , the stopped distribution is approximately (up to the fact that the process can be absorbed at zero) a binomial distribution. Summing up over individuals, we obtain the sum of binomially distributed random variables which is again a binomially distributed random variable. As the process has a negative expectation, the distributions are right-skewed and approach a normal distribution as t increases. Thus, for a decreasing joy of gambling, in the individual choice treatment, we would not expect a left-skewed distribution as is predicted for the contest.

Remark 3 (*Skewness of stopped value distribution*) *For the individual choice problem with subjects displaying decreasing joy of gambling, the model predicts that the distributions of the stopped values are not left-skewed.*

B Example Instructions

Moderate Negative, Contest-Lottery Order

Welcome

You are about to participate in a session on interactive decision-making. Thank you for agreeing to take part. The session should last about 90 minutes.

You should have already turned off all mobile phones, smart phones, mp3 players and all such devices by now. If not, please do so immediately. These devices must remain switched off throughout the session. Place them in your bag or on the floor besides you. Do not have them in your pocket or on the table in front of you. The entire session, including all interaction between you and other participants, will take place through the computer. You are not allowed to talk or to communicate with other participants in any other way during the session. You are asked to follow these rules throughout the session. Should you fail to do so, we will have to exclude you from this (and future) session(s) and you will not receive any compensation for this session.

We will start with a brief instruction period. Please read these instructions carefully. They are identical for all participants in this session with whom you will interact. If you have any questions about these instructions or at any other time during the experiment, then please raise your hand. One of the experimenters will come to answer your question.

Structure of the session

There are three parts to this session. Instructions for all parts are detailed below. Part 1 consists of 5 periods, during which you will be given the opportunity to familiarise yourself with the environment and the interface that forms the basis of the tasks to be completed in parts 2 and 3. Part 2 consists of 10 periods. At the beginning of each period, you will be matched with another participant for a decision task. Part 3 consists of 10 periods of a similar task. However, it will be an individual decision task, and the calculation of your earnings for this part will be different. Consequently, parts 1, 2 and 3 are related, but independent tasks.

Compensation for participation in this session

You will be able to earn money for your decisions in parts 2 and 3 of the session. What you will earn from part 2 will depend on your decisions, *the decisions of others* and chance. What you will earn from part 3 will *only* depend on your decisions and chance. Further details will be given in due course.

In the instructions, and all decision tasks that follow, payoffs are reported in experimental currency units (ECU). At the end of the experiment, your earnings in ECUs will be converted to Euros at the rate of 100 ECU: 1 Euro. Your final payment will be 3 Euros plus the sum of your earnings from parts 2 and 3. Final payment takes place in cash at the end of the session. Your decisions and earnings in this session will remain anonymous.

Instructions for Part 1

1. Part 1 consists of 5 periods—periods 1 through 5 of the experiment as a whole. During these periods you will *not* interact with other participants and your choices will *not* affect your earnings for this experiment. Part 1 is meant as practice to give you an opportunity to familiarise yourself with the environment and the computer interface.
2. During each period you will see a graph evolve in real-time. This graph will be referred to as a *random process* and has the following features:
 - (a) The process starts at a value of 15.
 - (b) You have 15 seconds before the value starts to randomly fluctuate.
 - (c) During this random fluctuation, every quarter of a second there is a 47% chance that the value increases by 1 and a 53% chance it decreases by 1. Consequently, if the value does not increase by 1 it decreases by 1, and vice-versa.
 - (d) This random fluctuation continues until the *time limit* is reached, which is 90 seconds, or until the value becomes zero. If the value becomes zero the process remains zero for the remainder of the period.
3. Your task is to select a value by “stopping” the process. The value when you stop the process is referred to as the *stopped value* for the period.
4. To do this, you have three controls at your disposal:

- (a) A *stop now* button: Clicking on this button will stop the process immediately.
- (b) An *upper bound*: If the current value of the process is equal to (or greater than) this number, the process will stop. By default this value is set to 45. You can change the value by clicking on the up or down arrows next to the number or by dragging the upper of the two horizontal blue lines on the graph, which represents this upper bound.
- (c) A *lower bound*: If the current value of the process is equal to (or less than) this number, the process will stop. By default this value is set to 0. You can change the value by clicking on the up or down arrows next to the number or by dragging the lower of the two horizontal blue lines on the graph, which represents this lower bound.

5. The stopped value is determined as follows:

- The current value of the process when you click on the “stop now” button.
Note: The stop now button *is active* during the 15 seconds before the random fluctuation begins. That is, pressing the button during this time will give you a stopped value equal to the start value.

- The current value of the process should it be greater than or equal to the current value of your upper bound.

Note: The upper bound *is not active* during the 15 seconds before the random fluctuation begins. That is, while you can change the value of the upper bound during this time, it will not stop the process until the random fluctuation has begun. As soon as this time is over, the control becomes active. In particular, if the upper bound is *equal to or below* the start value when the random fluctuation begins, you will stop the process immediately with a stopped value equal to the start value.

- The current value of the process should it be less than or equal to the current value of your lower bound.

Note: The lower bound *is not active* during the 15 seconds before the random fluctuation begins. That is, while you can change the value of the lower bound during this time, it will not stop the process until the random fluctuation has begun. As soon as this time is over, the control becomes active. In particu-

lar, if the lower bound is *equal to or above* the start value when the random fluctuation begins, you will stop the process immediately with a stopped value equal to the start value.

- The current value of the process at the time limit if none of the above applies before this time.
6. Stopping your process will not stop the evolution of the graph on your computer screen. It fixes your stopped value for the current period, which will be displayed on the right-hand side once it has been determined. Once you have stopped your process, you cannot change your stopped value for the current period, and your controls (the stop now button, the upper bound and the lower bound) will no longer be active for the remainder of the period.
 7. After the time limit has passed, you will have 15 seconds before the next period begins.

Example Random Processes and Control Questions

8. Before we begin with part 1, you will be shown some example realisation of the random process and asked some control questions to ensure you have understood the basic ideas of the environment.
9. The examples will be displayed on the computer screen. The examples are the same for all participants that you might interact with.
10. After the examples, you will be asked some control questions. Please read each question carefully and fill in what you think is the correct answer. If you do not get the correct answer for a question—the computer interface will prevent you from proceeding to the next question if this is the case—and cannot work out what the correct answer should be, please raise your hand and an experimenter will come to your cubicle to assist you.

Instructions for Part 2

11. Part 2 consists of 10 periods—periods 6 through 15 of the experiment as a whole. At the beginning of a period, you will be paired with another participant. We refer to

this participant as your match. This pairing procedure is repeated at the beginning of every period in part 2. The pairing is anonymous—you will not know the identity of your match, and vice-versa—and is determined randomly.

12. During a period you will control a random process similar to that described in part 1.
 1. The process will have the same features and controls as in part 1.
13. At the same time your match will also control a random process. The random process has the same features as your random process, as described above. However, it will be a different realisation. That is, your and your match's processes will be independent of each other. You will only ever see your own process and your decisions have no effect on the random fluctuation of your match's process, and vice-versa.
14. **Your payoff** for this period is determined by comparing your stopped value with that of your match's. If your stopped value is strictly larger than your match's, you will earn 150 ECU for this period. If your match's stopped value is strictly larger than yours, you will earn 0 ECU for this period. If the stopped values are exactly the same, then the computer will randomly determine who earns 150 ECU and who earns 0 ECU, with you or your match having an equal chance to be the one that earns 150 ECU.
15. Once both your and your match's process has stopped fluctuating, you will be shown your stopped value, your match's stopped value, your payoff for the period, and your match's payoff for the period.
16. Note that, if the stopped values were exactly the same, you will find out at this point who was randomly chosen to earn 150 ECU, with the other earning 0 ECU.
17. After a period is finished, you will be paired for a new period. Part 2 consists of 10 such periods.
18. Your payoff for part 2 is the sum of your payoffs from all the periods in part 2. That is, the sum of your payoffs from periods 6 through 15.

Instructions for Part 3

1. Part 3 consists of 10 periods—periods 16 through 25 of the experiment as a whole. Unlike part 2, you will not be paired with another participant during part 3. That is, what you will earn from part 3 will only depend on your decisions and chance.
2. During a period you will control a random process similar to that described in part 1.
 1. The process will have the same features and controls as in part 1.
3. **Your payoff** for this period is determined by a lottery. The outcome of this lottery can be either 150 ECU or 0 ECU. The chance of earning 150 ECU is determined by your stopped value for the period. In particular, the chance you will earn 150 ECU is equal to $(\text{stopped value} / 30) \times 100\%$, if your stopped value is less than 30; otherwise it is 100%.
4. Thus, your chance of earning the 150 ECU is increasing in your stopped value. If you have a stopped value of zero, you will earn 0 ECU for sure. If you have a stopped value of 15, you have a 50% chance of earning 150 ECU, since $(15/30) \times 100\% = 50\%$. If you have a stopped value of 30 or more, you will get 150 ECU for sure.
5. Once your process has stopped fluctuating, you will be shown your stopped value, your chance of earning 150 ECU from your lottery and your payoff for the period—that is, the outcome of your lottery for the period.
6. After a period is finished, you will move on to the next period. Part 3 consists of 10 such periods.
7. Your payoff from part 3 is the sum of your payoffs from all periods in part 3. That is, the sum of your payoffs from periods 16 through 25.
8. After part 3 has finished, the session will move on to a final questionnaire. An experimenter will come to your cubicle to start the questionnaire. Instructions for the questionnaire will be displayed on your computer terminal. Please read them carefully and proceed at your own pace.

C Screenshots

Figure C.1: Screenshot: Warm-up Phase

Figure C.2: Screenshot: Active Phase

Figure C.3: Screenshot: Feedback Phase

D Further Results

D.1 Additional Tables

Table D.1: Summary of the Outcomes in the Contest: Comparison across Data Subsamples

	Percent Stopped		Stopped Time (sec)		Stopped Value (ECU)	
Treatment	Time > 0		Average (95% C.I.)	Predicted	Average (95% C.I.)	Predicted
All Rounds						
Min -ve	94.0	***	23.77 (19.16, 28.38)	30.66	14.22 (13.67, 14.77)	14.54
Mod -ve	82.9	***	9.41 (7.49, 11.33)	10.75	12.93 (12.46, 13.40)	13.71
Ext -ve	35.1	***	0.50 (0.29, 0.70)	0.72	14.40 (14.18, 14.62)	14.53
First Five Round						
Min -ve	92.5	***	23.80 (19.88, 27.71)	30.66	13.76 (12.98, 14.53)	14.54
Mod -ve	83.8	***	10.04 (7.56, 12.53)	10.75	13.00 (12.39, 13.61)	13.71
Ext -ve	44.0	***	0.75 (0.38, 1.12)	0.72	14.10 (13.71, 14.49)	14.53
Last Five Rounds						
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: ***1%, **5%, *10% significance.

Table D.2: Summary of the Outcomes in the Lottery: Comparison across Data Subsamples

	Percent Stopped		Stopped Time (sec)		Stopped Value (ECU)	
Treatment	Time > 0		Average (95% C.I.)	Predicted	Average (95% C.I.)	Predicted
All Rounds						
Min -ve	94.1	***	21.90 (19.07, 24.72)	0.00	14.31 (13.98, 14.64)	15.00
Mod -ve	79.5	***	8.97 (6.99, 10.94)	0.00	12.36 (11.86, 12.85)	15.00
Ext -ve	44.2	***	0.64 (0.42, 0.85)	0.00	14.23 (14.01, 14.44)	15.00
First Five Round						
Min -ve	93.5	***	21.17 (17.36, 24.98)	0.00	14.38 (13.73, 15.02)	15.00
Mod -ve	79.8	***	9.44 (6.97, 11.92)	0.00	11.99 (11.18, 12.80)	15.00
Ext -ve	52.2	***	0.66 (0.43, 0.90)	0.00	14.10 (13.79, 14.40)	15.00
Last Five Rounds						
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: ***1%, **5%, *10% significance.

Table D.3: p -Values from Point-Prediction Tests for Hypotheses 1–3

Treatment	Contest		Lottery	
	H 1c	H 2b	H 3	H 3
	Stopped Time	Stopped Value	Stopped Time	Stopped Value
Min -ve	0.016	0.785	0.000	0.041
Mod -ve	0.055	0.001	0.000	0.000
Ext -ve	0.000	0.027	0.004	0.000

Notes: p -values are based on a Wald test of the estimated coefficients of a linear random-effects regression (with clustering at the matching group level) of either the stopped time or stopped value on a full set of treatment indicators. All tests use data from the last five rounds of the contest and the lottery parts, respectively.

Table D.4: Random-Effects Regressions of Behaviour across Periods: Probability of Gambling, Stopped Time and Stopped Value.

	Min -ve		Mod -ve		Ext -ve	
<i>Gambles</i>						
Contest	0.02	(0.188)	0.08***	(0.004)	−0.26***	(0.001)
Pay Order	0.01	(0.623)	0.10**	(0.042)	−0.15**	(0.027)
Contest × Pay 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)
<i>Stopped Time</i>						
Contest	6.59**	(0.015)	0.41	(0.778)	−0.95**	(0.025)
Pay Order	0.52	(0.861)	1.36	(0.219)	−0.95**	(0.025)
Contest × Pay Order	−9.45***	(0.009)	−0.08	(0.972)	1.79*	(0.066)
Period	0.14	(0.376)	−0.22**	(0.041)	−0.10**	(0.017)
<i>Stopped Value</i>						
Contest	0.58	(0.468)	0.32	(0.559)	0.66***	(0.000)
Pay Order	0.09	(0.881)	−0.79	(0.139)	0.43*	(0.060)
Contest × Pay Order	−1.33	(0.300)	0.49	(0.689)	−0.98***	(0.002)
Period	0.05	(0.285)	0.04	(0.448)	0.09***	(0.000)

Notes: The regression model is a random-effects probit model for the gambles dependent variable, and a linear random effects model for stopped time and stopped value dependent variables. Coefficients report the marginal effect on the dependent variable. Standard errors clustered at the matching group level. *** 1%, ** 5%, * 10%.

Table D.5: Complete Set of Contrast Tests from Regressions on Behaviour across Periods

Dependent Variable	Within-Session Tests (Phase 2 v Phase 3)						Between-Session Tests (Standard v Reverse Order)			
	Overall		Standard Order		Reverse Order		Phase 2		Phase 3	
	Period Trend		Cont v Lott		Cont v Lott		Cont v Lott		Cont v Lott	
	β	p	β	p	β	p	β	p	β	p
<i>Min -ve</i>										
Gambles	0.06	0.104	0.65	0.123	-0.61	0.163	0.41	0.209	-0.37	0.342
Stopped Time	0.14	0.376	6.59	0.015	-2.85	0.065	6.07	0.142	-2.33	0.445
Stopped Value	0.05	0.285	0.58	0.468	-0.75	0.186	0.49	0.387	-0.66	0.009
<i>Mod -ve</i>										
Gambles	-0.01	0.753	0.65	0.001	-0.35	0.130	-0.21	0.512	0.51	0.196
Stopped Time	-0.17	0.097	1.18	0.306	-0.28	0.813	-0.95	0.626	1.84	0.281
Stopped Value	0.04	0.448	0.32	0.559	0.82	0.282	1.11	0.000	0.03	0.959
<i>Ext -ve</i>										
Gambles	-0.11	0.000	-0.77	0.000	0.06	0.851	-0.33	0.000	-0.37	0.246
Stopped Time	-0.07	0.001	-0.62	0.000	0.34	0.114	-0.00	0.994	-0.28	0.228
Stopped Value	0.09	0.000	0.66	0.000	-0.31	0.133	0.23	0.308	0.12	0.342

Notes: The regression model is a random-effects probit model for the gambles dependent variable, and a linear random effects model for stopped time and stopped value dependent variables. β columns report the point prediction of the associated linear combination of coefficients from the (latent) linear regression. For the period trend and the standard-order contest-lottery contrast, this is just the coefficient associated with the period and contest variables, respectively. For the reverse-order contest-lottery contrast, this is contest plus contest x pay order. For the phase 2 contest-lottery contrast, this is contest minus pay order. For the phase 3 contest-lottery contrast, this is pay order plus contest plus contest x pay order. Standard errors are clustered at the matching group level. The p columns report the result of a two-sided test against the null hypothesis that the linear combination is equal to zero.

Table D.6: Aggregate-Level Stopping Choices across All Rounds

Treatment	Stopped Time > 0	Stopped with Button	Stopped with Threshold	Stopped with Bankruptcy	Never Stopped
<i>Contest</i>					
Min-ve	0.94	0.34	0.50	0.10	0.00
Mod-ve	0.83	0.29	0.49	0.05	0.00
Ext-ve	0.35	0.11	0.24	0.01	0.00
<i>Lottery</i>					
Min-ve	0.94	0.25	0.57	0.11	0.01
Mod-ve	0.80	0.24	0.51	0.05	0.00
Ext-ve	0.44	0.11	0.32	0.00	0.00

Notes: Numbers report the proportion of observations from each part. Data from rounds 1-10 for each part.

Table D.7: Subject-Level Averages: Characterising Stopping Choices across All Rounds

Treatment	Contest				Lottery			
	Never	Minority	Majority	Always	Never	Minority	Majority	Always
<i>How often did subject not stop immediately (i.e. gamble time > 0)</i>								
Min-ve	0.00	0.03	0.96	0.78	0.01	0.03	0.97	0.76
Mod-ve	0.03	0.12	0.84	0.54	0.03	0.14	0.80	0.51
Ext-ve	0.19	0.66	0.31	0.03	0.14	0.54	0.38	0.11
<i>How often did subject hit bankruptcy</i>								
Min-ve	0.46	1.00	0.00	0.00	0.44	0.99	0.00	0.00
Mod-ve	0.66	1.00	0.00	0.00	0.72	1.00	0.00	0.00
Ext-ve	0.94	1.00	0.00	0.00	0.95	1.00	0.00	0.00
<i>How often did subject stop using the stop button (strictly) during the active phase</i>								
Min-ve	0.25	0.59	0.31	0.00	0.36	0.74	0.16	0.00
Mod-ve	0.33	0.69	0.21	0.01	0.25	0.80	0.10	0.01
Ext-ve	0.57	0.93	0.05	0.00	0.68	0.91	0.09	0.01
<i>How often did subject stop using threshold interface (strictly) during the active phase</i>								
Min-ve	0.04	0.46	0.45	0.06	0.05	0.36	0.61	0.12
Mod-ve	0.10	0.50	0.42	0.07	0.05	0.44	0.49	0.05
Ext-ve	0.33	0.78	0.15	0.00	0.25	0.69	0.26	0.03

Notes: The proportions report the proportion of subjects whose stopping choices were never according to the stopping characteristic (in italic), accorded in a minority of rounds (no more than four out of ten), accorded in a majority of rounds (at least six out of ten), and accorded in all rounds. These proportions are reported separately for the contest and lottery parts. Data from round 1-10 for each part.

Table D.8: Heterogeneity in Response to change in Payoff Structure: Probability of Gambling across Periods

	Min -ve		Mod -ve		Ext -ve	
<i>Random-Effects model</i>						
Contest	0.02	(0.188)	0.08*** (0.004)		−0.26*** (0.001)	
Pay Order	0.01	(0.623)	0.10** (0.042)		−0.15** (0.027)	
Contest × Pay 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)	
<i>Random-Coefficients model</i>						
Contest	0.02	(0.214)	0.08*** (0.005)		−0.28*** (0.001)	
Pay Order	0.00	(0.664)	0.08* (0.078)		−0.16** (0.022)	
Contest × Pay Order	−0.02	(0.325)	−0.08** (0.027)		0.28 (0.125)	
Period	0.00	(0.339)	−0.00 (0.749)		−0.04*** (0.000)	

Notes: Random-effects probit models. The second regression is the same as the first one except that it includes a random coefficient for the contest variable and its interaction term with pay order. Coefficients report the marginal effect on the dependent variable. Standard errors clustered at the matching group level. *** 1%, ** 5%, * 10%. The likelihood Ratio test for significance of the random-coefficients model, compared to the random-effects model, gives $p < 0.01$ for all three drift parameters.

Table D.9: Further Statistics on the Distribution of Stopped Values

Task Order	Contest % of Stopped Values					Lottery % of Stopped Values				
	Mean	Median	< 15	= 15	> 15	Mean	Median	< 15	= 15	> 15
<i>Min -ve</i>										
Contest-Lottery	14.58	14	51	6	43	14.53	15	44	9	47
Lottery-Contest	13.87	15	44	12	44	14.09	15	46	11	43
All	14.22	15	47	9	44	14.31	15	45	10	45
<i>Mod -ve</i>										
Contest-Lottery	12.89	14	51	18	32	12.94	15	46	31	23
Lottery-Contest	12.97	15	47	24	30	11.77	13	59	16	25
All	12.93	15	49	20	31	12.36	14	52	24	24
<i>Ext -ve</i>										
Contest-Lottery	14.22	15	32	57	10	14.47	15	28	67	5
Lottery-Contest	14.58	15	18	76	6	13.99	15	46	46	8
All	14.40	15	25	66	8	14.23	15	37	57	6

Notes: Data from all rounds of the contest and lottery tasks.

Table D.10: Stopped Value Variance and Skewness in the Mod-ve Treatment

Task Order	Contest			Lottery		
	Variance	Skewness	p-Skewness	Variance	Skewness	p-Skewness
<i>Mod-ve: All Data</i>						
Contest-Lottery	41.02	-0.204	0.092	29.89	-0.675	0.000
Lottery-Contest	36.51	-0.517	0.000	40.47	0.010	0.933
All	38.71	-0.349	0.000	35.48	-0.298	0.001
<i>Mod-ve: Last Five Rounds</i>						
Contest-Lottery	39.48	-0.351	0.041	29.77	-0.772	0.000
Lottery-Contest	37.37	-0.431	0.013	39.31	-0.179	0.287
All	38.42	-0.392	0.002	34.56	-0.439	0.000
<i>Mod-ve: Period 15 versus Period 16</i>						
Contest-Lottery	36.69	-0.208	0.544	33.73	-0.882	0.019
Lottery-Contest	33.48	-0.920	0.015	46.34	-0.134	0.695
All	35.35	-0.543	0.042	39.67	-0.394	0.132

Notes: p-Skewness reports the p-value of a test for deviations from normality based on skewness (i.e. testing for deviation from zero skewness). For the lottery-contest order, subjects do not receive any feedback on the choices of other participants until after they have made their stopping decision for period 16.

D.2 Additional Figures

(a) Observed

(b) Predicted

Figure D.1: Behaviour across Periods by Task Order: Stopped Time

(a) Observed

(b) Predicted

Figure D.2: Behaviour across Periods by Task Order: Stopped Value

Figure D.3: Proportion of Rounds Subject Gambled (Stopped Value > 0) during the Last Five rounds.

Figure D.4: Proportion of Rounds Subject Gambled and Stopped using the Button Interface during the Last Five rounds.

Figure D.5: Proportion of Rounds Subject Gambled and Stopped using Threshold Interface during the Last Five rounds.

Figure D.6: Proportion of Rounds Subject Gambled and went Bankrupt during the Last Five rounds.

(a) All Rounds

(b) Last 5 Rounds

Figure D.7: Stopped Value Distribution in the Mod-ve Treatment