World Development 173 (2024) 106427

Contents lists available at ScienceDirect

World Development

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

Agricultural investment behaviour and contingency: Experimental evidence from Uganda

Ben D'Exelle [a,*], Alistair Munro [b], Arjan Verschoor [a]

[a] University of East Anglia, UK
[b] National Graduate Institute for Policy Studies (GRIPS), Japan

A R T I C L E I N F O

JEL classification:
C93
D03
D81
O13

Keywords:
Investment behaviour
Contingency
Decision-making under uncertainty
Uganda

A B S T R A C T

Underinvestment in agriculture – a major cause of rural poverty – may be due to difficulties in detecting 'contingency', defined as the influence one may exert on the outcome of a decision-making situation. Recently experienced contingency may create a mismatch between perceived and actual contingency in an investment decision-making situation, leading to sub-optimal investment behaviour. To test this, we use an experiment with poor farmers in Uganda used to low levels of contingency, as many factors (e.g., the weather, pests, price fluctuations) obscure the link between farm investment and outcomes. We find that in situations in which some contingency is present, investment levels respond positively to recently experienced contingency. In situations in which no contingency is present ('non-contingency'), investment responds negatively to recently experienced non-contingency. The findings that perceived contingency influences investment behaviour, and perceived contingency can be readily changed, may inform new behavioural policies to promote agricultural investment.

1. Introduction

Underinvestment in agriculture is held responsible for a major part of the persistence of poverty in the rural areas of developing countries (World Bank, 2007). Previously identified factors responsible for such underinvestment include risk aversion (Klasing, 2014), credit and insurance market failures, lack of access to information and resources (e.g., labour at peak times), paucity of savings instruments, and the unavailability of good agricultural extension services (see Foster and Rosenzweig (2010) and Wiggins et al. (2021) for reviews). We add another factor: the difficulty to detect 'contingency'. Contingency is a concept we borrow from psychology, in which it is defined as the extent to which one can influence the outcomes of a situation (Abramson et al., 1978).

When decision-makers have no influence over the outcomes of a situation, that situation is said to exhibit non-contingency. Conversely, the more influence a decision-maker has over the outcomes of a situation, the higher is the degree of contingency of that particular decision-making situation. Thus, contingency is a feature of a decision-making situation and may vary in degree. Applied to an investment situation, non-contingency implies that the investor may exert no influence on the distribution of returns on investment. The higher the degree of contingency of an investment situation is, the more influence on the distribution of returns on investment may be exerted.

While contingency and non-contingency are objective features of a decision-making situation, the amount of contingency that a decision-maker experiences (or perceived contingency) may be different from the amount that is objectively present. The harder it is to detect the relationship between cause and effect in a particular decision-making situation, the lower will be the amount of contingency that a decision-maker tends to experience. This is particularly pertinent for the population of interest in this paper: smallholder farmers in the Global South.

Smallholder farmers in the Global South frequently experience low contingency. They are used to situations in which their own influence on investment returns may be difficult to detect, as many factors may obscure the link between farm investment (e.g., new crops, or the use of inputs such as fertilizer, improved seeds, etc.) and outcomes such as crop yields or farm profits: the soil, the weather, plant pests and disease, price fluctuations, among many other factors (Dercon, 2008; Fafchamps, 2003). In addition, farming in the Global South is heavily affected by climate change (Mendelsohn, 2008), and in many locations, the impact of climate change on local weather patterns is disrupting established knowledge of cause and effect for smallholders.

Importantly, contingency may vary in degree and farmers may perceive more or less contingency than actually exists, often influenced

* Corresponding author.
E-mail addresses: b.dexelle@uea.ac.uk (B. D'Exelle), alistair-munro@grips.ac.jp (A. Munro), a.verschoor@uea.ac.uk (A. Verschoor).

https://doi.org/10.1016/j.worlddev.2023.106427
Accepted 2 October 2023
Available online 12 October 2023
0305-750X/© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

by the recent experience of contingency. In this study, we use an experiment to test the hypothesis that the amount of contingency recently experienced affects the amount of contingency perceived in a current investment situation, which in turn influences investment decisions.[1] This is an important question to answer for the following two reasons. First, it might confirm that a mismatch between the contingency of a recent situation and a present one may lead to sub-optimal investment behaviour. For example, the recent experience of non-contingency may lead investors to underestimate contingency when facing new investment opportunities, which leads to underinvestment. Second, if we confirm that perceived contingency can be readily changed, it would open new ways to promote investment behaviour. Finding ways of communicating contingency effectively may then need to be part and parcel of policies or programs that aim to promote agricultural investment.

To investigate the effect of experienced contingency and non-contingency on subsequent investment behaviour, we use an experiment with a sample of poor farmers in Uganda. In the experiment, our approach is to measure the effect of contingency on behaviour in an indirect way. We do so by first experimentally varying exposure to it with one game, followed by two more games that create settings where either some contingency or non-contingency exist. If we find that behaviour in the last two games correlates with prior exposure to contingency in the expected way, then we interpret this as evidence in support of an effect of experienced contingency. Specifically, in stage 1 of the experiment, subjects play a card guessing game in which we randomly vary between subjects the exposure to contingency and non-contingency. In stage 2, all subjects play an investment game, in which we vary contingency in two within-subject conditions. Investment success depends on correctly detecting a weight difference between two objects. In the first condition, the weight difference may be detected, while in the second, that is not humanly possible. The first condition is thus characterized by more contingency than the second, which is characterized by non-contingency. In stage 3, all subjects are invited to solve a puzzle in which success is impossible for the untrained player, but it takes costly time for subjects to realize this (subjects pay for each second played): a task with non-obvious non-contingency.

Combining stages 1 and 3 allows us to test the hypothesis that more recently experienced non-contingency leads to lower investment in a situation characterized by non-contingency. Combining stages 1 and 2 allows us to test the hypothesis that prior exposure to (a) contingency increases investment where contingency is present, and (b) to non-contingency decreases investment where non-contingency is present.[2]

To derive these hypotheses, we develop a theoretical model that relies on the following two assumptions. First, we assume that recently experienced (non-)contingency influences perceived (non-)contingency through a 'recency' heuristic.[3] A plausible motivation for the use of such a heuristic is that fully exploring contingency in each new situation is inefficient, and that the recent past is the best available approximation of the situation currently faced (Denrell & March, 2001; Teodorescu & Erev, 2014). Second, perceived (non-)contingency translates into a subjective probability that influences the level of risk-taking (investment levels).[4]

We confirm our hypotheses for both investment stages of the experiment, for stage 2 and stage 3. First, we find that more contingency experienced in stage 1 increases investment in stage 2 for the situations with more contingency. From this finding we infer that more contingency recently experienced increases the positive influence that subjects believe they may exert on the probability of a successful outcome. Second, those who experience non-contingency in stage 1 invest less costly time in solving the puzzle in stage 3 (which is characterized by non-contingency). This shows that a mismatch between the contingency of a recent situation and a present one may lead to sub-optimal investment behaviour: subjects who did not recently experience non-contingency waste more valuable resources since they are slower to detect non-contingency in the new situation. We also detect important heterogeneity in both effects, along locus of control and risk preferences.

Our study contributes to two strands of literature in development economics. First, development economists have increasingly used insights from behavioural economics (see e.g., Datta and Mullainathan (2014), Kremer et al. (2019), McKenzie et al. (2022)), to study decision-making and behaviour in many domains. One domain that has received much attention is risk-taking by poor farmers (see e.g., Verschoor and D'Exelle (2022) and references therein). To this literature, we contribute insights on the role of perceived contingency. Previous research has found effects on risky choice of an internal locus of control (Pinger et al., 2018; Salamanca et al., 2020).[5] People with a greater internal locus of control ('internals') have a greater tendency to attribute outcomes to their own influence rather than to external factors. For a given actual contingency and other things equal, internals will therefore perceive more contingency than externals will. Importantly, while locus of control is a personality characteristic that cannot be easily changed (Cobb-Clark & Schurer, 2013), perceived contingency is more malleable as it also depends on recently experienced contingency (Abramson et al., 1978), which we demonstrate in this paper. We also find an interesting interaction between both, with a stronger internal locus of control increasing the degree to which one is influenced by the recent experience of (non-)contingency.

Second, many studies have looked at technology adoption by farmers and the role of learning, through individual experimenting, or information from experimenting by other farmers, often supported by extension services (see e.g., Bandiera and Rasul (2006), BenYishay and Mobarak (2019), Carter et al. (2021), Conley and Udry (2010), D'Exelle and Verschoor (2023), Krishnan and Patnam (2013), Maertens (2017), Vasilaky and Leonard (2018)). We add the insight that in the noisy decision-making environments that farmers face, contingency may not always be detected, which might lower farmers' willingness to explore new investment opportunities. Farmers may believe that allocating resources to exploring the potential influence on the likelihood of successful investment outcomes is wasteful when too often the same conclusion is reached (that little or no such potential influence exists).

[1] A well-established research insight in psychology is that the amount of currently perceived contingency is influenced by the amount of recently experienced contingency (Cohen et al., 1976; Overmier, 1996). In typical experiments on contingency, some subjects first encounter a problem that cannot be solved through their own efforts, such as an impossible puzzle, e.g. Cohen et al. (1976). Exposure to the impossible problem lowers persistence and problem solving ability in subsequent trials even when the subsequent problems can actually be solved. Many of the initial experiments involved physical irritants such as dissonant noises (Hiroto, 1974), but later experiments have usually dispensed with the painful tasks and typically involve puzzles, word problems and similar challenges.

[2] We use a lab setting to test these hypotheses, as the issue is complex: using an RCT to exogenously vary exposure to (non-)contingency and study its effect in settings that vary in the actual presence of (non-)contingency may be prohibitively difficult.

[3] On the use of heuristics in decision-making under uncertainty see Charness and Levin (2005), Dohmen et al. (2009), Plous (1989) and Tversky and Kahneman (1974).

[4] Similar to previous literature on subjective probabilities, inspired by Tversky and Kahneman (1974), where cognitive heuristics are assumed to be a source of bias, we assume that a recency heuristic biases subjective probabilities towards recent experiences. Specifically, if the exploration of the features of a decision-making problem relies on the recent experience of (non-)contingency, then the assessment of probabilities will be biased by that experience.

[5] Other studies have found a link with illusion of control (Fellner, 2009; Thompson, 1999), confidence (Merkle, 2017; Murad et al., 2016), grit (Alaoui & Fons-Rosen, 2021), and optimism (Dohmen et al., 2018).

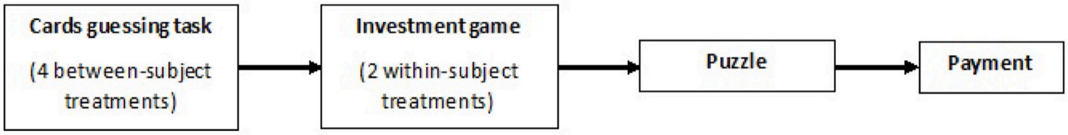

Fig. 1. Sequence of tasks.

Table 1
Treatments of cards guessing task.

Treatments	Rule	
T1: Straightforward rule	Card with the bird is correct	Contingent
T2: Complex rule	Card is correct if it has a red letter at the top; if neither card has a red letter at the top, the card with a yellow letter is correct	Non-contingent
T3: Confusing rule	Any card with a letter is correct	Non-contingent
T4: No task	No cards used	–

If such passivity becomes normal, ways would need to be found to help them recognize that exploring investment opportunities is worth doing. Our finding that perceived contingency in investment situations influences investment behaviour, and that perceived contingency can be readily changed among farmers is therefore encouraging. It suggests that influencing perceived contingency is in certain situations a promising way to encourage investment.

An important caveat is that these insights are inferred from a lab-in-the-field experiment. Whereas we observe a causal link between experienced contingency and investment behaviour in that experiment, we are not in a position to observe a causal link with real-world investment behaviour. For that, we might have had to implement a randomized controlled trial (RCT). It is hard to see how an RCT would have allowed us to exogenously vary the contingency present in investment situations, although perhaps it is possible, and it would certainly make for interesting research. However, we would argue that our approach is also of interest. Indeed, the use of lab-in-the-field experiments in development economics is by now well established, and has been found to have good external validity so long as the behaviour of interest is adequately captured in the experiment (Verschoor et al., 2016).

The remainder of the paper is laid out as follows. In Section 2, we present our design, which includes a presentation of the different experimental tasks, treatments, hypotheses and the sample. Section 3 contains the main results, while Section 4 analyses the heterogeneity of the treatment effects. Section 5 concludes.

2. Experiment

Fig. 1 presents the sequence of the different tasks. To investigate the behavioural influence of prior experience of contingency or non-contingency, participants start with a cards guessing task, which manipulates exposure to contingency via four between-subject treatments. This stage is inspired by the body of research in psychology that experimentally varies exposure to contingency.[6] Thereafter, all subjects make two decisions in an investment game, once when the success probability is 50% and labelled as outside their control (the non-contingency condition), and once when it is well above 50% if careful discernment is applied (the contingency condition). If we find that prior experience to contingency or non-contingency affects the decisions, then we infer that the subjective success probability is affected by alertness to contingency. In a final task, all participants try to disentangle a metal puzzle, in which each second of effort is costly and there is a financial reward for successfully completing the puzzle. After these three tasks, their payment is determined. Only then will they know the outcome of their decision in the investment game. In the rest of this section, we present more details on each of the tasks and treatments.

2.1. Cards guessing task

To manipulate exposure to (non-)contingency, we use a card guessing task. Subjects are presented with ten pairs of cards, one pair at a time, and are asked to guess which card is correct, according to a rule that we do not share with them. So subjects are told that there is a rule according to which one card is right and the other wrong, but we do not tell them what that rule is. Rather, after each guess, subjects are told whether they guessed correctly, but no further feedback is given. The pairs of cards shown to subjects are depicted in Figure A.1 in online Appendix A.

Using this task, we implement four treatments (see Table 1) with each participant randomly assigned to only one of the treatments. Randomization was done at the individual level. These treatments vary in the rule that is used to determine which card is right. In treatment T1, the rule is "the card with the bird is correct". This means that if subjects are shown the first card and point to the card on the left, then they will be told that they identified the correct card. After the feedback, the second pair will then be shown. This rule was designed to bring about exposure to contingency.

Two other rules were designed to bring about exposure to non-contingency, in different ways. In both treatments T2 and T3, the rule was construed such that subjects would be right or wrong for no apparent reason related to their mental exertions. In treatment T2, that is the case because the rule is complex; in treatment T3, that is the case because the rule is trivial, since in effect, there are no wrong answers. In other words, in treatment T2 we expose subjects to non-contingency through complexity and in treatment T3 through triviality. It should be noted that we have no expectation about the differential impact of T2 and T3. Based on our reading of the relevant psychology literature (referred to above), we considered that non-contingency may result both from problems that are too difficult to solve and from problems that are too easy to solve. We decided to implement the two alternative sources of non-contingency as something that is empirically interesting (whether complexity or triviality matters more) without having clear priors on which matters more. Treatment T4 has no card-guessing task, which we use as control treatment. A comparison between this control treatment and the treatments that expose subjects to contingency or non-contingency allows us to identify the effect of the experience of contingency and non-contingency, respectively.

2.2. Investment game

The card guessing task is followed immediately by an investment game. In the first stage of the game, participants are asked to identify

[6] See Overmier (1996, 2002) for references to many relevant experimental studies in psychology.

the heavier of two small containers, weighing about 35 g each. After weighing the containers, they are asked to decide how many counters out of 20 to invest, each worth 400 Ugandan shillings. The investment is tripled if the correct container is identified to be the heavier one, and lost if not. The containers are transparent, and filled with locally popular nuts (so-called ground nuts).

The investment task is done twice, under two different conditions, the order of which is randomized at the individual level. The key difference between the two conditions is that in one condition, the contingency condition, the heavier container is 2 g heavier than the lighter one, whereas in the non-contingency condition, the weight difference is 0.2 g. According to evidence from psycho-physics, the latter difference cannot be detected, whereas many people, with effort, can detect the former difference in weight for a base weight of about 35 g (Gescheider, 1997, Ch. 1). We confirm this in our experiment (see the next section). Importantly, we tell subjects these facts: both the difference in weights in the two tasks and the fact that a 0.2 g difference is all but non-detectable, whereas a 2 g can be sensed by many people. Specifically, when subjects face the 2 g condition they are told (the experimental instructions are in Appendix H):

> *'We will show you later that one is heavier, using the machine that can detect small weight differences. The difference is small, but we know that if you try hard enough, many people can tell the difference. We really believe that not only machines can tell the difference in weight but many people can too, if they try hard.'*

We refer here to a 'machine', which in the actual experiment is a 'digital balance'. To this we add the (truthful) information that:

> *'We know this because we have tried this with hundreds of people before and two out of three people correctly decide which container is heavier.'*

On the other hand, with the 0.2 g condition the subjects are told:

> *'We really believe that it is impossible to tell the difference in weight between these two containers; the difference is too small for a person to feel it with their own hands. We will show you later, using the machine that can detect very small weight differences, that one is heavier but whether or not you get it right is a matter of luck. Why do we say that it is a matter of luck? Because we really believe that machines can tell the difference in weight, but people cannot.'*
>
> To this we add the (truthful) information that:
>
> *'We believe this because we have tried this with hundreds of people before and they guess the wrong container as often as they guess the right one.''*[7]

We then give participants exactly 30 s to weigh the containers before they must come to a decision.

2.3. Puzzle

In the third task, subjects are given a small, metal puzzle (see Figure A.2 in online Appendix A for examples) and told that they will win a prize if they can separate its parts. They are endowed with 3000 shillings and told that for each second of play they will lose 10 shillings from that endowment. On the other hand, if they solve the puzzle, then as a reward they will receive 5 times their remaining endowment. Subjects can stop playing at any moment, in which case they receive their remaining endowment. If they use up the whole endowment, then the task ends automatically. Conversely, they can simply refuse to do the task in which case they receive the 3000 shillings of the original endowment.

Prior to setting the terms of the challenge, the experimenter demonstrates that the puzzle can be solved by disentangling it. They do not explain how this is achieved and do the disentangling at some distance from the participant. In practice, solving the puzzle can usually only be achieved through consulting instructions. The task is one of non-obvious non-contingency, since it is much more complex than it appears at first. Indeed, no subject managed to disentangle the puzzle. After two checks for understanding, the participant is asked if they wish to do the task. If they do, then a stopwatch is used to determine the time at which they stop. The outcome and time is recorded and subjects are paid once all other subjects in the session have completed the task.[8]

2.4. Treatment predictions

To guide our analysis, we develop treatment predictions for the investment game and the puzzle, using some simple theory.[9]

Each of the decisions in the investment game is a version of the standard portfolio problem. This is a classical problem in the economics of uncertainty, with many real-life applications (Gollier, 2001, Ch. 4). In the standard portfolio task, an agent decides how much of a sure wealth to invest in a risk-free asset and in a risky asset. In our experiment, subjects' initial endowment is equal to E, $E-\alpha$ of which is invested in the risk-free asset and α in the risky asset. The return of the risk-free asset is zero in our experiments, so that the portion $E-\alpha$ of the endowment not invested in the risky asset is simply retained. The return of the risky asset is a random variable $\tilde{y}$. The value of the portfolio is equal to $E-\alpha+\alpha\tilde{y} = E+\alpha(\tilde{y}-1)$, which the agent decides on by choosing α, subject to the constraint $0 \leq \alpha \leq E$. The expected value of $(\tilde{y}-1)$ is set to be greater than zero in our experiment or risk-averse decision makers would not invest.

In our experiment, the probability distribution of $\tilde{y}$ is a two-point distribution: with probability of success p the investment is multiplied by a factor $k > 1$ and with probability $(1-p)$, the investment is lost. Subjects thus face the prospect $[p,\ E+\alpha(k-1);(1-p),\ E-\alpha]$, in which α is freely chosen subject to the constraint mentioned. However, subjects believe that they face prospect $[p^*,\ E+\alpha(k-1);(1-p^*),\ E-\alpha]$, since the success probability is not unambiguously known. Experiments are designed to create conditions that we expect to cause the perceived probability p^* of success to deviate from p. The way we test for discrepancies between p and p^* is, in a between-subject design, to manipulate conditions for p^* in various ways while holding p constant. If investment responds to treatment, then we infer that p^* is sometimes a biased estimate of p.

To show this formally, let us use a utility function that has the constant relative risk aversion (CRRA) form, $u(y) = y^{1-\gamma}/(1-\gamma)$ for some risk aversion parameter γ. Using first order conditions of utility maximization in an expected utility framework, utility would be maximized if α is chosen such that:[10]

$$p^*(k-1)(E+\alpha(k-1))^{-\gamma} = \left(1-p^*\right)(E-\alpha)^{-\gamma} \tag{1}$$

[7] We used the opportunity provided by other experiments to try out these weighing tasks in other villages in Uganda.

[8] After respondents complete all stages – but before their payment is calculated – we also elicit a measure of confidence that the chosen container is correct. We do this to gain some insight into whether the subjects actually viewed the two tasks as significantly different. At the same time, we made it unincentivized to avoid raising the complexity of the experiment. For details see the instructions in Appendix H.

[9] Note that the design means that objectively, expected payoffs are linearly increasing in the investment in the weighing tasks and linearly decreasing in time spent on the puzzle, so conclusions and predictions about investment and time spent on the puzzle carry over to conclusions and predictions about expected payoffs.

[10] When $\gamma = 1$ utility takes the log form and $\alpha/E = (kp^*-1)/(k-1)$. In this case, α/E is also clearly increasing in p^*. For $\gamma < 0$, agents will invest all their endowment.

With $\mu = \left(\frac{p^*(k-1)}{1-p^*}\right)^{-1/\gamma}$, this can be rewritten as:

$$\frac{\alpha}{E} = \frac{(1-\mu)}{(k-1)\mu + 1} \tag{2}$$

For risk averse agents ($\gamma > 0$), the effect of an increase in p^* is positive:

$$\frac{\partial \alpha / E}{\partial p^*} = \frac{\mu k}{p^*(1-p^*)\gamma\,((k-1)\mu+1)} \tag{3}$$

An increase (decrease) in p^* as a result of prior experience of contingency (non-contingency) would then increase (decrease) the α at which utility is maximized. In the rest of this section, we use this to derive predictions about the treatment differences.

If subjects believe the message about the weight differences, then investment levels should be higher in the 2 g condition. In addition, if recent experience of (non-)contingency influences alertness to (non-)contingency – through a 'recency' heuristic (see references used in the introduction) – then the investment made in the 2 g condition should be higher in treatment T1 compared to treatment T4, and the investment made in the 0.2 g condition should be lower in treatments T2 and T3 compared to treatment T4.[11] This translates into our first hypothesis.

Hypothesis 1. Prior exposure to contingency increases investment where contingency is present (2 g condition), while prior exposure to non-contingency decreases investment where non-contingency is present (0.2 g condition).

For the puzzle task, effort is positively related to the player persisting with trying to solve the puzzle. All players are shown that the puzzle is theoretically solvable, but it is much more complex than it appears at first. The task is thus one of non-obvious non-contingency. Participants differ in their prior exposure to non-contingency in the cards guessing game, which we expect to influence their effort in the puzzle game. We set out a brief, formal model in Appendix C, but the intuition is the same as for the investment game and so is the resulting hypothesis:

Hypothesis 2. Prior exposure to non-contingency (treatments T2 and T3) decreases investment in stage 3, i.e. the costly time spent on the puzzle.

2.5. Implementation and sample

The site chosen for the experiment was Sironko and Lower Bulambuli, which together are the former Sironko district in eastern Uganda. Approximately 300,000 people live in the selected region of around 1270 square kilometers, 95% of whom are small-scale farmers (Verschoor et al., 2016). Within the site, we randomly selected 20 villages. For each selected village we compiled a list of all adult (over 18) village members and from that list randomly selected up to 20 adults (with no more than one person per household). After the experiment we conducted a brief livelihoods survey with each participant.

Overall, we have 395 subjects. Both assignment to the four card guessing treatments (T1 to T4) in stage 1 and task order for the investment games in stage 2 were randomized at the individual level. Subjects were guided individually along the various tasks by trained enumerators. Several experimental rooms were used, in each of which the materials required for the tasks were laid out. Only one subject was present in each room at a time, to ensure confidentiality of decision making. From start to finish (including the livelihoods survey), each subject spent no longer than 45 min participating in the research. Average payout was 7450 shillings, or almost two daily agricultural wages of 4000 shillings. Nearly all subjects answered correctly the control questions about the investment game (96.71%) and the puzzle (94.18%).[12] Table B.1 in the online Appendix provides descriptive statistics of relevant socio-economic characteristics of the participants by treatment. We observe that the four card treatments are balanced across all characteristics.

In the card guessing task, the mean number of times the correct card was identified is 7.98 in treatment T1 and 5.88 in treatment T2. In treatment T1 (T2) the mean number of right answers in the final five questions was 0.41 higher (0.46 lower) than in the first five and 81 (49) subjects out of 98 (99) performed no worse in the last half of the card guessing task. Thus performance in treatment T1 was significantly higher than in treatment T2 (two-sided $p = 0.000$, Mann–Whitney test) and the gain between the first and second half was also significantly greater (two-sided $p = 0.000$, Mann–Whitney test). By the end of the task, 75 subjects in treatment T1 claimed to know the correct rule, while only 5 in treatment T2 were confident enough to make the same claim. On this basis, we claim that treatments T1 and T2 were effective at influencing the experience of contingency and non-contingency. For treatment T3 everyone answered all card-guessing questions correctly by construction, so we cannot make the same claims for that treatment.

3. Results

3.1. Investment game

Before we compare investment levels across treatments and conditions, we need to verify that (1) in the 2 g condition the correct container was selected more often than in the 0.2 g condition, and (2) this is the case with all treatments. The first condition tests the assumption of increased contingency in the 2 g condition, while the second should hold as the treatments only influence 'perceived' contingency and not 'actual' contingency.

The percentage that identified the heavier container correctly for the 0.2 g difference was 48.5%, which is not significantly different from 50% (two-sided $p = 0.546$, binomial test), whereas at 60.3% the percentage of correct guessers for the 2 g difference was significantly different from 0.5 (two-sided $p = 0.000$, binomial test). Fig. 2 shows the breakdown by treatment. In all treatments, the fraction choosing correctly in the 2 g case was higher than the corresponding fraction for the 0.2 g case. For treatments T3 and T4 the difference is statistically significant at the 5% level (two-sided $p = 0.005$ and 0.039 respectively, McNemar's test). Differences across treatments but within container type are not significantly different at the 5% level. For the 2 g containers, the percentage of accurate guesses was significantly above 50% for all but treatment T2 whereas in the 0.2 g case, in all treatments the difference from 50% was not significant.

In Fig. 3 we observe that the number of counters invested is higher with 2 g than with 0.2 g, and this is true for all the cards treatments. Using a paired t-test, we find that within-participant differences in investment between the 2 g and 0.2 g conditions are positive and statistically different from zero, for each of the four treatments (two-sided $p < 0.03$). So when contingency is actually present (2 g condition),

[11] It is worth noting that the expected channel of influence is through contingency being perceived, not greater accuracy in detecting weight differences. Prior research in psychology is equivocal about whether we should expect higher accuracy in treatment T1 compared to treatments T2, T3 and T4. Eisenberger (1992) notes that many early experiments do not produce evidence of 'learned industriousness' (e.g., Maier & Seligman, 1976), but when there is no effective ceiling on performance (e.g., Eisenberger et al., 1976) priming contingency is associated with higher effort and performance. In our case, we cap the evaluation time for the task at 30 s and, anyway, it is not obvious that further time would improve subjects' accuracy. Thus we have no strong prior that treatment T1 should produce higher accuracy than treatments T2, T3 and T4. What we expect instead is that, for given accuracy in detecting a weight difference, a clearer perception of contingency makes it more likely that the detected weight difference translates into higher investment.

[12] See online Appendix H for the control questions we used.

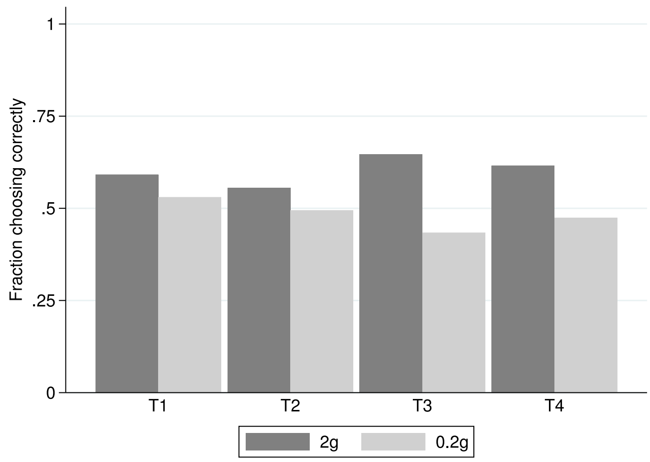

Fig. 2. Fraction choosing correctly, by treatments.

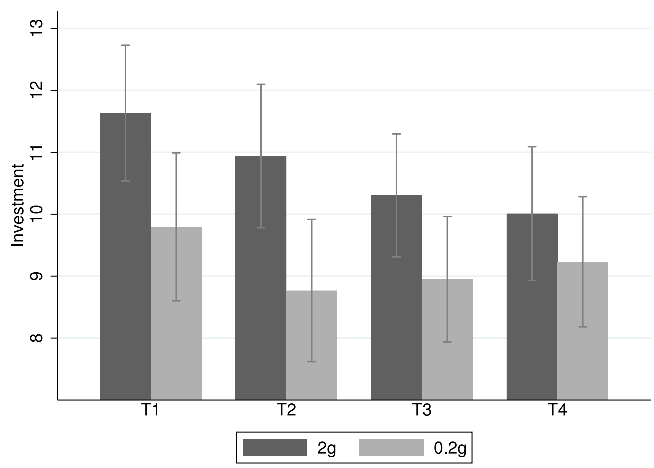

Fig. 3. Investment levels, by treatments (95% confidence intervals).

subjects on average are discerning: they respond to the increased success probability. Prior exposure to non-contingency (in T2 and T3) has not undone that. We summarize this in our first finding.

Finding 1. *People's investment behaviour responds to actual contingency and is not undone by any treatment.*

In Fig. 3 we also observe that with the 2 g condition, investment increases from T4 (control) up to T1 (prior exposure to contingency). In contrast, we do not observe such a pattern with the 0.2 g condition. To analyse this in more detail we run a regression, with the number of coins invested as dependent variable and the treatment categories as main explanatory variable. We use treatment T4 (control) as reference category. Table 2 presents the results.[13] In Models 2, 4 and 6 we control for gender, age, education, land, household size, task order, locus of control and risk preferences of the participant.[14] In all regressions we also control for the order of the 2 g and 0.2 g conditions, and we use experimenter fixed effects. Standard errors are estimated with bootstrapping clustered at the village level.[15]

Table 2
Investment: Treatment differences.

	(1) 2 g	(2) 2 g	(3) 0.2 g	(4) 0.2 g	(5) 2 g–0.2 g	(6) 2 g–0.2 g
T1	1.611***	1.674***	0.538	0.592	1.073**	1.082**
	(0.508)	(0.417)	(0.478)	(0.460)	(0.460)	(0.435)
T2	0.923	0.908	−0.477	−0.508	1.400***	1.416***
	(0.757)	(0.827)	(0.530)	(0.638)	(0.456)	(0.430)
T3	0.356	0.535	−0.163	0.031	0.519	0.503
	(0.722)	(0.665)	(0.787)	(0.725)	(0.438)	(0.403)
Constant	11.001***	7.476***	11.221***	6.395***	−0.220	1.081
	(0.828)	(1.378)	(0.868)	(1.749)	(0.457)	(0.918)
R^2	0.035	0.146	0.109	0.263	0.089	0.137
Controls	No	Yes	No	Yes	No	Yes

Notes: N = 395. OLS regression. ***, **, * indicate two-sided significance levels at 1, 5, and 10%, respectively. Standard errors (in parentheses) estimated with bootstrapping (2000 iterations) clustered at village level. Experimenter fixed effects are used in all models. In Models 2, 4 and 6 we also control for gender, age, education, land, household size, task order, locus of control and risk preferences of the participant. See Table D.1 in the online Appendix for the coefficients of the controls. The results are robust to excluding the participants who did not answer all control questions correctly, as demonstrated by Table E.1 in the online Appendix. For results with treatments T2 and T3 pooled see Table F.1 in the online Appendix.

The coefficient of T1 (recent experience of contingency) is statistically significant with the 2 g condition (columns 1–2), while the coefficients of T2 (non-contingency, complex) and T3 (non-contingency, trivial) are not. With the 0.2 g condition (columns 3–4) none of the coefficients is statistically significant. Columns 5 and 6 use the within-difference between the conditions 2 g and 0.2 g as the dependent variable. This removes potential biases that affect both conditions equally. For example, if one of the between-subject treatments not only influences recent experience of (non-)contingency, but also triggers other factors (e.g., confidence, practice, or trust in the experimenters), the within-difference filters these factors out. We observe that the coefficient of T2 is now also statistically significant.[16] We summarize these observations in a second finding.

Finding 2. *In the contingency condition (2 g), prior exposure to contingency (T1) increases investment. In the non-contingency condition (0.2 g), investment is not influenced by prior exposure to non-contingency. Recent experience of contingency (T1) or non-contingency (T2) increases the difference in investment levels between both conditions.*

[13] Table B.2 shows that a large proportion of participants invest exactly 50% of the counters. Investing 50% of the counters is compatible with a range of values of risk aversion, though it might also be a focal point for some people. This does not raise any problems for the regression analysis. A visual inspection of the residuals shows that they are normally distributed, which is also confirmed by a skewness and kurtosis test for normality (p-values are 0.300 and 0.348 for the residuals of columns 2 and 4). As none of the p-values are below 10%, the Gaussian assumptions of the linear regression model are not rejected.

[14] To measure risk preferences we used a hypothetical question, adapted from Dohmen et al. (2005), about subjects' willingness to invest $x \in \{0, 20000, 40000, 60000, 80000, 100000\}$, so up to 100,000 shillings in an asset that yields a return of 100 percent if successful and minus 50 percent if a failure, with equal probability. 100,000 shillings is about a month's worth of daily wages for an agricultural worker. Subjects chose one of six decision cards on which the two outcomes of a possible choice were clearly displayed. 0: refers to 0 invested; 5: refers to 100,000 invested. Thus this variable is equal to zero for respondents with the highest risk aversion, and higher values indicate lower risk aversion.

[15] Here we follow Cameron et al. (2008) who advise to use bootstrapping if the number of clusters is low. We have 20 clusters as participants were recruited from 20 villages.

[16] We also analysed whether the treatments influence the participants' confidence that they chose the correct container. The results are reported in Appendix G. The effects go in the expected direction, but lack statistical significance, which is probably due to the elicitation being done at the end of the experiment (after all tasks were completed) and being unincentivized.

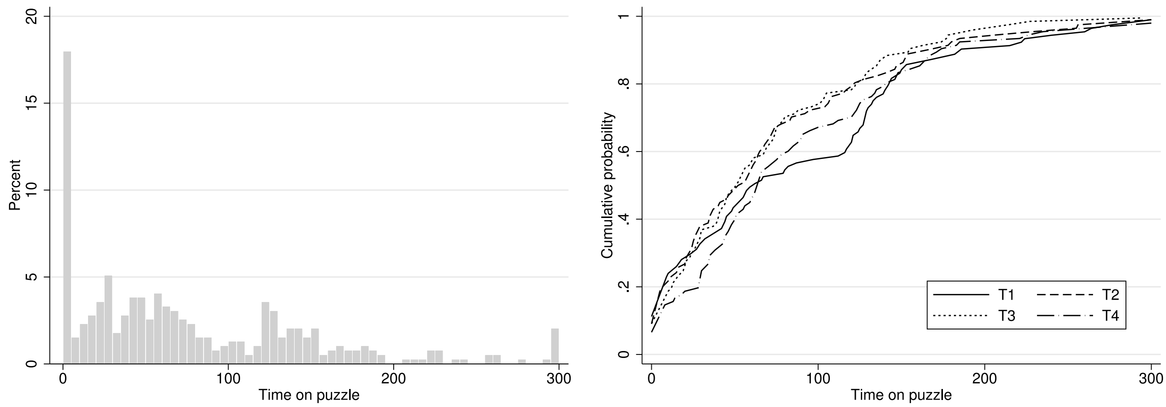

(a) Distribution of time (b) Time by treatments

Fig. 4. Time on the puzzle.

3.2. Puzzle

Panel (a) of Fig. 4 summarizes the overall distribution of time spent on the puzzle. Nearly 17.97% of subjects (71) declined the invitation to try the puzzle. Their times are recorded as zero. The distribution of stopping times thereafter is quite uneven, but eight subjects continued right until the end and therefore took home nothing from the experiment, as no participant managed to solve the puzzle. This confirms that there is no contingency in the puzzle task. Panel (b) breaks down the cumulative distribution of stopping times by treatment and suggests some differences for intermediate stopping times. In particular, in the T2 and T3 treatments more participants tend to spend less time on the puzzle than in treatments T1 and T4.

To test whether these treatment differences are statistically significant we use a regression. Since approximately 18% of all subjects do not try the puzzle at all, we use a tobit specification. Table 3 presents the marginal effects on the likelihood that the puzzle is tried and the time spent on the puzzle conditional on trying the puzzle. We observe that prior exposure to contingency does not have any significant effect in any of the models. Prior exposure to non-contingency (T2 and T3), in contrast, has a negative influence both on participants' propensity to try the puzzle and on the time they spend on the puzzle. The effects are sizeable: they are around 6 percentage points less likely to try the puzzle and if they decided to try the puzzle they spent more than 10 s less on it. We summarize these results in a new finding

Table 3
Puzzle: Treatment differences.

	(1) Try = 1	(2) Seconds | Try = 1	(3) Try = 1	(4) Seconds | Try = 1
T1	−0.008	−1.639	−0.003	−0.530
	(0.057)	(11.083)	(0.054)	(10.596)
T2	−0.059*	−10.168*	−0.063**	−10.699*
	(0.031)	(5.648)	(0.032)	(5.590)
T3	−0.079**	−13.122**	−0.064*	−10.931*
	(0.031)	(5.619)	(0.036)	(6.456)
Controls	No	No	Yes	Yes

Notes: N = 395. Tobit regression. ***, **, * indicate two-sided significance levels at 1, 5, and 10%, respectively. Models 1 and 3: marginal effect on Try = 1; models 2 and 4: marginal effect on seconds | Try = 1. Standard errors (in parentheses) estimated with bootstrapping (2000 iterations) clustered at village level. Experimenter fixed effects are used in all models. In Models 3 and 4 we also control for gender, age, education, land, household size, locus of control and risk preferences of the participant. See Table D.2 in the online Appendix for the coefficients of the controls. The results are robust to excluding the participants who did not answer all control questions correctly, as demonstrated by Table E.2 in the online Appendix. For results with treatments T2 and T3 pooled see Table F.2 in the online Appendix.

Finding 3. *Prior exposure to non-contingency decreases the likelihood of trying the puzzle and the time spent on the puzzle.*

4. Heterogeneity: Locus of control and risk preferences

In this section, we investigate whether the treatment effects vary along individual risk preferences and locus of control. Both would plausibly interact with perceived contingency: risk preferences since the perception of contingency may itself be thought of as guided by a probability distribution; locus of control since, for a given amount of contingency objectively present in a situation, internals would tend to perceive more contingency than externals.

4.1. Locus of control

Locus of control is a psychological trait, which refers to the factors that people tend to attribute outcomes to. People with an internal locus of control tend to attribute outcomes of situations to their own influence, and people with an external locus of control to outside factors. We hypothesize that having an internal locus of control increases the response to (non-)contingency of a situation. Put differently, people with an external locus of control attribute outcomes to outside factors, and would therefore be less able to discern whether a situation has actual (non-)contingency. In a similar way, we expect that a stronger internal locus of control increases the degree to which one can be influenced by recent experience of (non-)contingency. In the theoretical models we developed, this would translate in a stronger change in p^{*}, and hence a stronger effect on behaviour.

To test this hypothesis, we add an interaction term between the treatment indicator variables and a measure of external locus of control (LOC). We then estimate the marginal effects at three levels of LOC, determined by the midpoint of three intervals of similar frequency. A lower level of LOC corresponds with a lower level of external locus of control, and hence a higher level of internal locus of control. Table 4 presents the results.

The first two columns report the treatment differences on investment behaviour, separately for each of the three LOC levels, and the 2 g and 0.2 g conditions. In Column 1, we observe that the effect of recent experience of contingency (T1) on investment is only statistically significant among the participants with the lowest levels of LOC, i.e., participants with the highest internal locus of control. The coefficients, however, are not statistically different between the lowest and highest LOC levels. None of the coefficients of the effects of recent experience of non-contingency (T2 and T3) are statistically different from zero. Also, no effects are identified where there is no contingency (Column 2).

Table 4
Heterogeneity by locus of control.

		(1) 2g		(2) 0.2g		(3) Try = 1		(4) Seconds | Try = 1	
T1	LOC = 1	1.954***	(0.483)	0.712	(0.518)	−0.047	(0.053)	−8.737	(9.756)
	LOC = 2	1.709***	(0.427)	0.602	(0.472)	−0.009	(0.054)	−1.790	(10.450)
	LOC = 3	0.975	(1.068)	0.272	(0.878)	0.091	(0.070)	20.223	(15.162)
	LOC(1) − LOC(3)[a]	0.65		0.21		7.85***		7.93***	
T2	LOC = 1	1.145	(0.781)	−0.181	(0.647)	−0.071**	(0.034)	−12.666**	(6.148)
	LOC = 2	0.924	(0.814)	−0.491	(0.634)	−0.064*	(0.034)	−11.009*	(6.035)
	LOC = 3	0.259	(1.469)	−1.423	(1.051)	−0.040	(0.061)	−6.266	(9.754)
	LOC(1) − LOC(3)[a]	0.44	(0.507)	1.54	(0.215)	0.29		0.47	
T3	LOC = 1	0.743	(0.718)	0.298	(0.758)	−0.091***	(0.035)	−15.537**	(6.258)
	LOC = 2	0.543	(0.645)	0.017	(0.713)	−0.063	(0.039)	−10.865	(6.958)
	LOC = 3	−0.057	(0.998)	−0.824	(0.937)	0.019	(0.067)	3.380	(11.692)
	LOC(1) − LOC(3)[a]	0.54		1.66		5.10**		4.33**	

Notes: N = 395. Columns 1 and 2: OLS regression, marginal effects reported. Column 3: Tobit regression, marginal effect on Try = 1; Column 4: Tobit regression, marginal effect on seconds conditional on Try = 1. Standard errors (in parentheses) estimated with bootstrapping (2000 iterations) clustered at village level. Experimenter fixed effects are used in all models. We also control for gender, age, education, land, household size, locus of control and risk preferences of the participant. ***, **, * indicate two-sided significance levels at 1, 5, and 10%, respectively.
[a] Chi-square test with H_0 being that the marginal effects are the same for LOC = 1 and LOC = 3.

Columns 3 and 4 test whether locus of control interacts with the treatment effects on whether the puzzle is tried and the time spent on the puzzle conditional on trying it. In the T1 panel, we find that while the coefficients differ significantly between the highest and lowest locus of control, none of the coefficients are statistically different from zero. In the T2 and T3 panels, the coefficients are statistically significant among respondents with the highest level of internal locus of control. The sign of the interactions is as expected: respondents with higher internal locus of control are less likely to try the puzzle and spend less time on it conditional on trying it. In panel T3, the difference between the highest and lowest locus of control is also statistically significant, as demonstrated by the significant chi-square test. These results suggest that the previously observed negative effect of the prior experience of non-contingency is driven by participants with the highest internal locus of control. This is in line with what we hypothesized. We summarize these results in a new finding.

Finding 4. *The positive effect of recent experience of contingency on investment in the contingency condition is driven by participants with an internal locus of control. The negative effect of recent experience of non-contingency on the likelihood of trying the puzzle and the time spent on the puzzle, is driven by participants with an internal locus of control.*

4.2. Risk preferences

As demonstrated in Section 2.4 above and in Appendix C, both the optimal investment and the time spent on the puzzle decrease with the risk aversion parameter γ. This is consistent with the significant coefficients of individual risk preferences in Tables D.1 and D.2, which report the coefficients of the controls used. Apart from a direct effect, individual risk preferences could also moderate the effect of recent experience of (non-)contingency, which we will study next.

Starting with the investment game, according to the theory, the effect of changing p^* is stronger with less risk averse agents, i.e., lower γ. Eq. (4) shows the partial derivative of $\frac{\partial\alpha/E}{\partial p^*}$ with respect to γ for $\gamma \neq 1$ and an interior solution. For a risk averse person ($0 < \mu < 1$), this expression is negative, which implies that more risk averse agents are less sensitive to changes in the perceived probability of success, p^*.

$$\frac{\partial^2\alpha/E}{\partial\gamma\partial p^*} = -\frac{\partial\alpha/E}{\partial p^*}\frac{\mu(k-1)+1-ln(\mu)}{\gamma\,((k-1)\mu+1)} \tag{4}$$

Looking at the puzzle task, we know from the theoretical model presented in Appendix C that very risk averse people ($\gamma > 1$) will never try the puzzle, and a change in p^* will not change that. For lower γ, an increase in p^* increases the time spent on the puzzle, i, and the effect of an increase in p^* on i is stronger with smaller γ. In other words, the effect of a change in p^* will be stronger with stronger risk preferences.

To investigate the interaction with risk preferences, it should be noted that we cannot observe risk preferences in the investment game, since we do not directly observe p^*.[17] That is why we use a measure of risk preferences elicited with the livelihoods survey, in which probabilities are given and unambiguous. In particular, we use a hypothetical question, adapted from (Dohmen et al., 2005), about subjects' willingness to invest $x \in \{0, 20000, 40000, 60000, 80000, 100000\}$, so up to 100,000 shillings in an asset that yields a return of 100 percent if successful and minus 50 percent if a failure, with equal probability. 100,000 shillings is about a month's worth of daily wages for an agricultural worker. Subjects chose one of six decision cards on which the two outcomes of a possible choice were clearly displayed. 0: refers to 0 invested; 5: refers to 100,000 invested. Thus this variable is equal to zero for respondents with the highest risk aversion, and higher values indicate lower risk aversion. Even though these preferences were elicited after the experiment, they are not influenced by the experimental treatments. A regression of the risk preferences on the treatment indicators confirms this. None of the coefficients are statistically significant at the 10 percent level and the p-value of an F-test of joint significance is equal to 0.929.[18]

To investigate whether and how risk preferences interact with the treatment effects, we interact the risk preference measure with the treatment indicator variables in the regressions. We then estimate the marginal effects at four levels of risk preferences. Of the lowest three risk preference levels, we use the mid-point, as these levels are less common in our sample. Table 5 reports the results. The first two columns report the treatment differences for each of the four risk preference levels, for the 2 g and 0.2 g conditions, respectively.

We observe that the positive effect of recent experience of contingency (T1) on investment where there is contingency (2 g condition) is stronger among participants with stronger risk preferences. Among participants with the highest risk aversion (risk = 1), no effect is observed. Based on a chi-square test, the coefficients are statistically different between respondents of the lowest and highest levels of risk preferences.

In the panels of T2 and T3, we observe that also the recent experience of non-contingency increases investment among the participants

[17] Subjects believe they face the prospect $[p^*, 20+2\alpha; 1-p^*, 20-\alpha]$, in which p^* is the subjective success probability, and α the number of counters invested. We infer from the investment response to the prior experience of contingency an increase in p^*, but we do not directly observe p^*.

[18] For the highest value, individuals may also be risk neutral or risk loving.

Table 5
Heterogeneity by risk preferences.

		(1) 2 g		(2) 0.2 g		(3) Try = 1		(4) Seconds | Try = 1	
T1	risk = 1	−0.658	(1.407)	−2.422	(1.471)	−0.130	(0.094)	−17.846	(11.921)
	risk = 2	1.397***	(0.451)	0.240	(0.527)	−0.014	(0.053)	−2.580	(10.056)
	risk = 3	2.424***	(0.582)	1.571***	(0.493)	0.029	(0.051)	6.696	(11.861)
	risk = 4	3.451***	(1.075)	2.903***	(0.903)	0.061	(0.053)	17.055	(15.883)
	risk(1) − risk(4)[a]	3.09*		5.85***		3.18*		3.40*	
T2	risk = 1	−1.188	(1.980)	−1.865	(1.572)	−0.133**	(0.067)	−18.169**	(9.049)
	risk = 2	0.628	(0.920)	−0.685	(0.704)	−0.074**	(0.033)	−12.208**	(5.736)
	risk = 3	1.536**	(0.740)	−0.096	(0.558)	−0.048	(0.031)	−8.785	(5.998)
	risk = 4	2.445**	(1.037)	0.494	(0.813)	−0.025	(0.039)	−5.057	(8.035)
	risk(1) − risk(4)[a]	2.00		1.30		1.66		1.04	
T3	risk = 1	−1.567	(1.938)	−0.969	(1.551)	−0.002	(0.057)	−0.352	(9.926)
	risk = 2	0.264	(0.763)	−0.092	(0.766)	−0.059*	(0.034)	−10.115*	(5.911)
	risk = 3	1.180**	(0.569)	0.346	(0.693)	−0.089**	(0.040)	−14.923**	(7.370)
	risk = 4	2.096**	(0.976)	0.784	(0.954)	−0.118**	(0.056)	−19.683*	(10.334)
	risk(1) − risk(4)[a]	1.89		0.72		1.66		1.38	

Notes: N = 395. Columns 1 and 2: OLS regression, marginal effects reported. Column 3: Tobit regression, marginal effect on Try = 1; Column 4: Tobit regression, marginal effect on seconds conditional on Try = 1. Standard errors (in parentheses) estimated with bootstrapping (2000 iterations) clustered at village level. Experimenter fixed effects are used in all models. We also control for gender, age, education, land, household size, locus of control and risk preferences of the participant.

[a] Chi-square test with H_0 being that the marginal effects are the same for risk = 1 and risk = 4.

with the strongest risk preferences. The observation that investment in the 2 g condition is also increased by prior experience of non-contingency (T2 and T3) among participants with strong risk preferences might be due to 'grit' (Alaoui & Fons-Rosen, 2021). Participants with stronger risk preferences might be more inclined to think that a challenging decision-making situation will eventually yield to their determination.

The effects on investment where there is no contingency (0.2 g condition) are presented in Column 2. We observe a positive effect of recent experience of contingency, but only among the participants with the strongest risk preferences. Based on a chi-square test, the coefficients are statistically different between the lowest and highest levels of risk preferences. This shows that even where there is no (or very little) contingency, the recent experience of contingency does influence investment behaviour among the respondents with strongest risk preferences. This suggests that an additional mechanism might be at work. The most plausible candidate is 'confidence'. As documented by Murad et al. (2016), confidence tends to be positively associated with a greater willingness to take risk. It might therefore be that participants who have stronger risk preferences might have higher levels of confidence, which is triggered by the 'straightforward rule' in the cards guessing game (treatment T1). Note that this finding does not contradict the hypothesized mechanism that works via a change in perceived contingency, as the effect of treatment T1 is stronger in the 2 g condition compared to the 0.2 g condition.

Columns 3 and 4 test whether risk preferences interact with the treatment effects on the likelihood that the puzzle is tried and the time spent on the puzzle conditional on trying it. We observe that while the effect of recent experience of contingency (T1) varies between different levels of risk preferences (the chi-square test is marginally significant), the effect is not statistically significant for any of the separate risk preference levels. The recent experience of non-contingency (T2 and T3) lowers the likelihood that the puzzle is tried and the time spent on the puzzle. While we find that some of the coefficients are statistically significant, the effect of recent experience of non-contingency does not vary between different levels of risk preferences (the chi-square test is not statistically significant). We summarize the results of this section in a new finding.

Finding 5. *The positive effect of prior experience of contingency on investment increases with risk preferences, in both 2 g and 0.2 g conditions. The negative effect of prior experience of non-contingency on the likelihood that the puzzle is tried and the time spent on the puzzle does not interact with risk preferences.*

5. Conclusion

We can now spell out the main things we have learned about the effects of perceived (non-)contingency on investment. In the investment game, the presence of contingency in the 2 g condition means that, compared to perceived non-contingency, perceived contingency translates at the margin into a greater success probability. We show that perceived contingency responds to recent experience of contingency: recent experience of contingency raises investment in situations in which the success probability responds to perceived contingency. The recent experience of non-contingency does not affect investment. Neither do we detect any effects of recent experience of (non-)contingency on investment where there is no contingency (0.2 g condition). This provides partial support for Hypothesis 1. With the puzzle, there is no contingency, as the puzzle cannot be disentangled (none of the participants managed to do so). We found that recent experience of non-contingency decreases the likelihood of trying the puzzle and the time spent on the puzzle conditional on trying the puzzle. This provides support for Hypothesis 2.

The results thus display an asymmetry that is in line with our expectations. Prior exposure to contingency increases investment in the second game, when contingency is present, whereas prior exposure to non-contingency decreases investment in the third game, where non-contingency is present. It appears that subjects who were exposed to contingency became more alert to it in a new situation in which that was present, realized that they could influence the success probability of investment by taking advantage of the contingency, and raised their investment accordingly. By contrast, subjects who were exposed to non-contingency in the card guessing game were quicker than others to give up in the third game, the puzzle, in which no contingency was present. In their case, it appears that their prior exposure to non-contingency made them more alert to it in a new situation characterized by non-contingency, realized that they could not influence the success probability (i.e. it was and would remain practically zero) and therefore lowered their investment (of costly time) accordingly.

We also detected interesting heterogeneity in the effects along locus of control and individual risk preferences. First, the positive effect of recent experience of contingency on investment under the 2 g (contingency) condition is driven by participants with an internal locus of control. The negative effect of recent experience of non-contingency on the likelihood of trying the puzzle and the time spent on the puzzle is also driven by participants with an internal locus of control. This suggests that those with an internal locus of control are more responsive to recent experience of (non-)contingency than others. Second, the positive effect of recent experience of contingency on investment is driven by participants with strongest risk preferences. Risk preferences do not interact with the treatment effects in the puzzle task.

Our findings may inform new behavioural policies to promote agricultural investment. Since factors such as unknown features of the soil, ill-understood agricultural inputs and unidentified pests and diseases co-determine yields, the role of one's own efforts in determining the success of an agricultural investment that requires careful application (e.g., a new seed variety) is easily obscured and the belief that no contingency is present (even when it is) may therefore take hold. Our experiments show that being alert to the possible presence of contingency – a recent experience of being in control of outcomes through careful discernment – promotes investment in situations characterized by non-obvious contingency. Finding ways of communicating contingency effectively may need to be part and parcel of promoting agricultural investment.

Beyond this, the perception of contingency is a crucial step underpinning effective decision-making in many important domains. As Fryer and Leenknecht (2023) point out, drawing the attention of students to the contingency between study and understanding is an important element in motivating self-learning by children in schools. Similarly, successful entrepreneurship in many spheres of business is enhanced by an ability to seek and understand contingency (Newman et al., 2019). Thus the lessons of the paper may be applied to a variety of policy issues where the effectiveness of incentives is sensitive to the perception of contingency by individual agents.

Encouragingly, our study shows that perceived contingency can be readily changed. Indeed, some important ways in which entrepreneurship may be made more effective is through updating beliefs about contingency. For instance, if a fellow farmer demonstrates the correct agronomic practices for a new seed, then those who observe it may realize there is contingency in this situation (e.g. appropriate spacing increases the likelihood of a good harvest). In other words, they may realize that there is more they can do to influence the success of the investment than they previously thought. The contribution of our study is to make this part of the decision-making process explicit, and show that it is influenced by recent experience.

CRediT authorship contribution statement

Ben D'Exelle: Conceptualization, Methodology, Writing, Data analysis, Funding acquisition. **Alistair Munro:** Conceptualization, Methodology, Writing, Data analysis, Funding acquisition. **Arjan Verschoor:** Conceptualization, Methodology, Writing, Data analysis, Funding acquisition.

Declaration of competing interest

The authors declare that they have no relevant or material financial interests that relate to the research described in the paper.

Data availability

Data will be made available on request.

Acknowledgements

The research documented in this paper was financed by ESRC-DFID, UK grant ES/J008893/1. Some of Alistair Munro's expenses were supported by JSPS, Japan KAKENHI Grant Number 25101002. Ethical clearance for the study was obtained from the Uganda National Council for Science and Technology (UNCST) with research registration number SS2806, and from the University of East Anglia. We thank Joshua Balungira and his team at The Field Lab for support with the fieldwork and Borja Perez Viana for research assistance. We thank participants at seminars at UEA, ISS (The Hague), University of Manchester, Jawaharlal Nehru University (New Delhi), the Institute of Dalit Studies (New Delhi), and SEEDEC Wageningen for their comments.

Online appendix

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.worlddev.2023.106427.

References

Abramson, L. Y., Seligman, M. E., & Teasdale, J. D. (1978). Learned helplessness in humans: Critique and reformulation. *Journal of Abnormal Psychology*, *87*(1), 49–74.
Alaoui, L., & Fons-Rosen, C. (2021). Know when to fold'em: The flip side of grit. *European Economic Review*, *136*, Article 103736.
Bandiera, O., & Rasul, I. (2006). Social networks and technology adoption in northern mozambique. *The Economic Journal*, *116*(514), 869–902.
BenYishay, A., & Mobarak, A. M. (2019). Social learning and incentives for experimentation and communication. *Review of Economic Studies*, *86*(3), 976–1009.
Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, *90*(3), 414–427.
Carter, M., Laajaj, R., & Yang, D. (2021). Subsidies and the African green revolution: Direct effects and social network spillovers of randomized input subsidies in mozambique. *American Economic Journal: Applied Economics*, *13*(2), 206–229.
Charness, G., & Levin, D. (2005). When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect. *American Economic Review*, *95*(4), 1300–1309.
Cobb-Clark, D. A., & Schurer, S. (2013). Two economists' musings on the stability of locus of control. *The Economic Journal*, *123*(570), F358–F400.
Cohen, S., Rothbart, M., & Phillips, S. (1976). Locus of control and the generality of learned helplessness in humans. *Journal of Personality and Social Psychology*, *34*(6), 1049–1056.
Conley, T. G., & Udry, C. R. (2010). Learning about a new technology: Pineapple in Ghana. *American Economic Review*, *100*(1), 35–69.
Datta, S., & Mullainathan, S. (2014). Behavioral design: A new approach to development policy. *Review of Income and Wealth*, *60*(1), 7–35.
Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*(5), 523–538.
Dercon, S. (2008). Fate and fear: Risk and its consequences in Africa. *Journal of African Economies*, *17*(S2), ii97–ii127.
D'Exelle, B., & Verschoor, A. (2023). Village networks and entrepreneurial farming in Uganda. *World Development*, *167*, Article 106241.
Dohmen, T., Falk, A., Huffman, D., Marklein, F., & Sunde, U. (2009). Biased probability judgment: Evidence of incidence and relationship to economic outcomes from a representative sample. *Journal of Economic Behaviour and Organization*, *72*(3), 903–915.
Dohmen, T., Falk, A., Huffman, D., Sunde, U., Shu, J., & Wagner, G. G. (2005). Individual risk attitudes: New evidence from a large, representative experimentally-validated survey. *IZA Bonn discussion paper No. 1730*.
Dohmen, T. J., Quercia, S., & Willrodt, J. (2018). Willingness to take risk: The role of risk conception and optimism.
Eisenberger, R. (1992). Learned industriousness. *Psychological Review*, *99*(2), 248–267.
Eisenberger, R., Park, D. C., & Frank, M. (1976). Learned industriousness and social reinforcement. *Journal of Personality and Social Psychology*, *33*(2), 227–232.
Fafchamps, M. (2003). *Rural poverty, risk and development*. Cheltenham, UK: Edward Elgar.
Fellner, G. (2009). Illusion of control as a source of poor diversification: Experimental evidence. *The Journal of Behavioral Finance*, *10*(1), 55–67.
Foster, A. D., & Rosenzweig, M. R. (2010). Microeconomics of technology adoption. *Annual Review of Economics*, *2*(1), 395–424.
Fryer, L. K., & Leenknecht, M. J. (2023). Toward an organising theoretical model for teacher clarity, feedback and self-efficacy in the classroom. *Educational Psychology Review*, *35*(3), 68.
Gescheider, G. A. (1997). *Psychophysics: The fundamentals*. Psychology Press.
Gollier, C. (2001). *Economics of time and risk*. MIT Press Cambridge Mass..
Hiroto, D. S. (1974). Locus of control and learned helplessness. *Journal of Experimental Psychology*, *102*(2), 187—193.
Klasing, M. J. (2014). Cultural change, risk-taking behavior and implications for economic development. *Journal of Development Economics*, *110*, 158–169.
Kremer, M., Rao, G., & Schilbach, F. (2019). Chapter 5 - behavioral development economics. In B. D. Bernheim, S. DellaVigna, & D. Laibson (Eds.), *Handbook of behavioral economics: applications and foundations 1*: *Vol. 2*, *Handbook of behavioral economics - foundations and applications 2* (pp. 345–458). North-Holland.
Krishnan, P., & Patnam, M. (2013). Neighbors and extension agents in ethiopia: Who matters more for technology adoption? *American Journal of Agricultural Economics*, *96*(1), 308–327.
Maertens, A. (2017). Who cares what others think (or do)? Social learning and social pressures in cotton farming in India. *American Journal of Agricultural Economics*, *99*(4), 988–1007.
Maier, S. F., & Seligman, M. E. (1976). Learned helplessness: theory and evidence. *Journal of Experimental Psychology: General*, *105*(1), 3—46.

McKenzie, D., Mohpal, A., & Yang, D. (2022). Aspirations and financial decisions: Experimental evidence from the Philippines. *Journal of Development Economics*, *156*, Article 102846.

Mendelsohn, R. (2008). The impact of climate change on agriculture in developing countries. *Journal of Natural Resources Policy Research*, *1*(1), 5–19.

Merkle, C. (2017). Financial overconfidence over time: Foresight, hindsight, and insight of investors. *Journal of Banking & Finance*, *84*, 68–87.

Murad, Z., Sefton, M., & Starmer, C. (2016). How do risk attitudes affect measured confidence? *Journal of Risk and Uncertainty*, *52*(1), 21–46.

Newman, A., Obschonka, M., Schwarz, S., Cohen, M., & Nielsen, I. (2019). Entrepreneurial self-efficacy: A systematic review of the literature on its theoretical foundations, measurement, antecedents, and outcomes, and an agenda for future research. *Journal of Vocational Behavior*, *110*, 403–419.

Overmier, J. B. (1996). Richard L. Solomon and learned helplessness. *Integrative Physiological and Behavioral Science*, *31*(4), 331–337.

Overmier, J. B. (2002). On learned helplessness. *Integrative Physiological & Behavioral Science*, *37*(1), 4–8.

Pinger, P., Schäfer, S., & Schumacher, H. (2018). Locus of control and consistent investment choices. *Journal of Behavioral and Experimental Economics*, *75*, 66–75.

Plous, S. (1989). Thinking the unthinkable: The effects of anchoring on likelihood estimates of nuclear war. *Journal of Applied Social Psychology*, *19*(1), 67–91.

Salamanca, N., de Grip, A., Fouarge, D., & Montizaan, R. (2020). Locus of control and investment in risky assets. *Journal of Economic Behaviour and Organization*, *177*, 548–568.

Teodorescu, K., & Erev, I. (2014). Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological Science*, *25*(10), 1816–1869.

Thompson, S. C. (1999). Illusions of control: How we overestimate our personal influence. *Current Directions in Psychological Science*, *8*(6), 187–190.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131.

Vasilaky, K. N., & Leonard, K. L. (2018). As good as the networks they keep? Improving outcomes through weak ties in rural Uganda. *Economic Development and Cultural Change*, *66*(4), 755–792.

Verschoor, A., & D'Exelle, B. (2022). Probability weighting for losses and for gains among smallholder farmers in Uganda. *Theory and Decision*, *92*, 223–258.

Verschoor, A., D'Exelle, B., & Perez-Viana, B. (2016). Lab and life: Does risky choice behaviour observed in experiments reflect that in the real world? *Journal of Economic Behaviour and Organization*, *128*, 134–148.

Wiggins, S., Glover, D., & Dorgan, A. (2021). Agricultural innovation for smallholders in sub-Saharan Africa.

World Bank (2007). *World development report 2008: agriculture for development*. Washington, D.C.: World Bank.