

TI 2026-062/I  
Tinbergen Institute Discussion Paper

# Off-Path Beliefs

*Duarte Gonçalves<sup>1</sup>*

*Terri Kneeland<sup>2</sup>*

*Andreas G. B. Ziegler<sup>3</sup>*

<sup>1</sup> University College London

<sup>2</sup> University of Calgary

<sup>3</sup> Vrije Universiteit Amsterdam and Tinbergen Institute

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and Vrije Universiteit Amsterdam.

Contact: [discussionpapers@tinbergen.nl](mailto:discussionpapers@tinbergen.nl)

More TI discussion papers can be downloaded at <https://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam  
Gustav Mahlerplein 117  
1082 MS Amsterdam  
The Netherlands  
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam  
Burg. Oudlaan 50  
3062 PA Rotterdam  
The Netherlands  
Tel.: +31(0)10 408 8900

# Off-Path Beliefs

Duarte Gonçalves<sup>\*</sup>   Terri Kneeland<sup>†</sup>   Andreas G. B. Ziegler<sup>‡</sup>

## Abstract

The off path plays a key role in theoretical models of strategic interaction, yet, empirically, little is known about choices and beliefs off path. Using a novel experimental paradigm where off-path behaviour is well identified and mistakes are rare, we document three findings. First, participants overestimate mistake probabilities and overreact to payoff differences. Second, although off-path moves strongly predict subsequent deviations, participants underinfer from them when forming beliefs about later choices. Third, underinference is strongest among participants assigning low probability to off-path moves. They are perfection believers: following an off-path move, they continue ruling out subsequent deviations.

**Keywords:** Game Theory; Belief Formation; Sequential Rationality; Quantal Response; Incentives; Costly Cognition.

**JEL Classifications:** C72, C92, D83, D84, D91.

---

<sup>\*</sup> Department of Economics, University College London; duarte.goncalves@ucl.ac.uk.

<sup>†</sup> Department of Economics, University of Calgary; t.kneeland@ucalgary.ca.

<sup>‡</sup> School of Business and Economics, VU Amsterdam; mail@andreasziegler.eu.

*This draft: 27 August 2026.*

# 1. Introduction

The notion of the off path is central to dynamic games. In standard solution concepts, the path of play that players expect to happen crucially depends on the reasoning about what would happen off path, that is, after histories of play that players do not expect to occur. As any textbook treatment of the matter discusses, beliefs about behaviour off path determine the path of play in the first place and theoretical predictions differ depending exactly on these beliefs, on how they are restricted, and on how off-path moves are interpreted. For instance, a standard assumption—closely tied to perfection and the principle of sequential rationality (Kreps, 2023, Appendix 12)—is that an off-path move can be treated as an “extremely unlikely ‘mistake’ that is unlikely to occur again” (Mas-Colell et al., 1995, p. 282). More than a theoretical curiosity, subtle off-path considerations routinely underlie models used to assess antitrust and competition policy (Nocke and Whinston, 2010; Fershtman and Pakes, 2012), environmental regulation (Ryan, 2012), or arms races and conflict deterrence (Baliga and Sjöström, 2004; Chassang and Padró-i-Miquel, 2010).

Despite the often implicit but pervasive prominence of assumptions regarding the off path in shaping our understanding of strategic interactions, predictions, and policies, discussions about the off path have remained mostly normative rather than positive. Whether or not perfection is “an unreasonable principle on which to hang your hat” (Kreps, 2023, p. 591), evidence about how people actually reason about the off path is mostly absent from theoretical discussions. This is not surprising: direct evidence about histories that are truly not supposed to be reached is, by definition, hard to obtain. Moreover, in settings in which deviations are documented—such as centipede, ultimatum, and related dynamic games (McKelvey and Palfrey, 1992; Andreoni and Blanchard, 2006; Palacios-Huerta and Volij, 2009; Levitt et al., 2011; Healy, 2024)—deviations are often far from infrequent and are typically conflated with (and explainable by) uncertainty about players’ motivations, such as altruism and fairness concerns, social norms, risk attitudes, reputation building, or ad-hoc ‘crazy types’. Only with better evidence on off-path behaviour and beliefs, in settings where motivations are transparent and mistakes are rare, can theoretical models aspire to encode positive predictions based on general principles.

This paper studies choices and beliefs on and off path in a simple laboratory experiment where off-path moves are both unlikely and believed to be unlikely. We design an experiment in which off-path behaviour is directly observed, mistakes are well defined, and beliefs about behaviour following on- and off-path moves can be elicited separately.

We find that off-path moves, whilst rare, are payoff-dependent and very informative about subsequent choices, being tied to a higher mistake propensity of some participants. The data suggest that behaviour off-path reflects a selection effect. Beliefs, in turn, exhibit selection neglect. Finally, participants who assign the lowest probability to off-path moves appear to suffer most strongly from selection neglect.

The design is deliberately simple. Participants in a choice experiment make simple monetary choices.

They first choose between two monetary amounts in one-shot problems and then face sequential problems in which they choose one of two binary choice problems and subsequently choose one of the two monetary amounts in the selected problem. In most problems, what choice is or is not a mistake is ambiguous. In contrast, in our setting, choosing a lower monetary amount is an unambiguous mistake. Similarly, in sequential problems, choosing the first-stage option with the lower maximum continuation payoff is an off-path move. In a separate belief experiment, participants observe the same decision environments and report incentivised beliefs about the choices made by participants in the choice experiment.

Our design solves a basic identification problem. First, in contrast to other dynamic-game experiments, motivations are transparent and off-path moves are well-identified. Our design deliberately shuts off other-regarding preferences and higher-order belief considerations and features simple decisions. Second, the same binary choice problem can appear as a one-shot problem, as a continuation problem following an on-path move, or as a continuation problem following an off-path move, allowing us to compare choices and beliefs at the same decision problem across different path histories, while varying the payoff cost of deviations.

We first establish that mistakes are rare and payoff-dependent and beliefs about mistakes overestimate their frequency and their dependence on payoff differences. Participants choose the lower payoff in one-shot binary choice problems only 2.37% of the time, and move off path in the first stage of sequential problems in 5.73% of the instances. Participants in the belief experiment expect mistakes to be rare but significantly overestimate their frequency. Then, leveraging the random variation in payoffs we introduce in our design, we find that both choices and beliefs respond to costs of mistakes. In line with prominent models of costly cognition and stochastic choice (Luce, 1959; McFadden, 1974; Fudenberg et al., 2015; Caplin and Dean, 2015; Matějka and McKay, 2015), mistakes are less frequent, and are believed to be less frequent, when the payoff cost of a mistake is larger, but beliefs are much more sensitive to payoffs than choices actually are. This holds both in one-shot problems, considering the difference in payoffs, and in sequential problems, where the cost of an off-path move is the difference in continuation payoffs between on and off path.

The main analysis compares choices and beliefs following on- and off-path moves. The binary choice problems in the second stage are substantially more likely to generate a suboptimal choice when they follow an off-path move compared to following an on-path move or in one-shot problems. The effect is large: suboptimal choices increase by more than 10 percentage points after off-path moves. Thus, off-path moves, whilst rare, are nevertheless informative events: one deviation from optimality does imply that a subsequent deviation is significantly more likely. Moreover, the effect is larger when the first-stage deviation entails a larger difference in continuation payoffs: the greater the cost of an off-path move, the more informative it is about the continuation play.

Beliefs partly reflect but substantially underestimate this path effect. Participants in the belief experiment assign higher probabilities to mistakes after off-path moves and lower probabilities after

on-path moves. They also expect more severe off-path deviations to be followed by more subsequent mistakes. However, beliefs move too little when conditioning on off-path events. The difference between choices and beliefs is especially striking because beliefs are more responsive than choices when considering, for instance, payoff variations. When it comes to conditioning on the path, participants move in the right direction but underinfer from the information contained in off-path moves.

We next show that this underinference is naturally understood as selection neglect. Participants differ in their propensity to make mistakes: those who make more one-shot mistakes are also more likely to move off path and more likely to make suboptimal continuation choices. A simple participant-level logit model of payoff-dependent mistakes predicts the same selection pattern: participants with higher estimated mistake propensities are disproportionately selected into off-path histories. We then adapt a Grether-style belief-updating regression to quantify how beliefs respond to the empirical informativeness of path moves (Grether, 1980, 1992). Participants' beliefs reflect an understanding that path moves are informative, but substantially underuse this information: they incorporate only about 1/9 of the information contained in observing off-path moves. In our context, underinference in beliefs amounts to a failure to incorporate selection effects, which is closely related to other instances of selection neglect (Jehiel, 2018; Esponda and Vespa, 2018; Barron et al., 2024) and failures of contingent reasoning (Martinez-Marquina et al., 2019; Niederle and Vespa, 2023; Esponda and Vespa, 2024).

We also examine participants who assign zero probability to off-path moves, and separately those who assign close to zero probability to such moves. More than one in eight of the participants in the belief experiment assign probability zero to off-path moves in every first-stage belief report; we call them perfection believers. We find that these participants, in line with perfection arguments, do treat off-path moves as uninformative mistakes: they do not expect more continuation mistakes after off-path moves than on path. That is, even when asked to condition on a move that their own first-stage beliefs treated as impossible, they continue to report beliefs close to a perfect-behaviour model. In contrast, participants who assign a higher probability to off-path moves are also the ones whose beliefs more correctly anticipate the selection effects present in the choice data.

Lastly, response times provide auxiliary evidence that the payoff effects we document are related to cognitive frictions. Choices are slower when the relevant payoff comparisons are closer, especially in sequential problems, where choices are more sensitive to payoff differences. Belief reports display similar patterns: participants take longer when the payoff comparison makes others' behaviour less straightforward to assess. In addition, participants who spend more time reporting beliefs tend to report higher beliefs about others' mistake propensities, consistently with the idea that perfect behaviour is an instinctive benchmark, whereas thinking through when and why others make mistakes requires more deliberation.

## Related Literature

This paper contributes to a long literature, theoretical and experimental, on backward induction, subgame perfection, and dynamic games. It is motivated by a large theoretical literature on perfection, sequential equilibrium, and refinements on dynamic games (Selten, 1975; Kreps and Wilson, 1982; Cho and Kreps, 1987; Battigalli and Siniscalchi, 2002), as well as work emphasising the conceptual difficulty of rational behaviour in extensive-form games after unexpected histories (Binmore, 1987; Basu, 1990; Reny, 1992, 1995). Experimental work on centipede games documents systematic deviations from backward induction (McKelvey and Palfrey, 1992; Palacios-Huerta and Volij, 2009; Levitt et al., 2011). Other experiments show that small payoff changes, fairness concerns, or social preferences can overturn standard predictions in dynamic and static games (Goeree and Holt, 2001; Andreoni and Blanchard, 2006; Fehr and Schmidt, 1999). Closest to our focus on off-path reasoning, Wang (2023) studies beliefs and higher-order beliefs in centipede games, and Breitmoser et al. (2014), Lin and Palfrey (2024), and Fong et al. (2025) study off-path behaviour through the lens of quantal response and level- $k$  reasoning. Recent work elicits utilities and beliefs to examine where behaviour deviates from standard predictions in strategic settings (Healy, 2024).

Our contribution is empirical and our design complementary to existing experimental work. Rather than studying equilibrium play in a rich strategic game, we remove strategic-preference confounds and use simple monetary choices to identify mistakes directly. This allows us to separate trembles from selection and to compare beliefs following on- and off-path moves at the same continuation problem and directly assess reasoning about off-path moves of others.

A second related literature elicits beliefs in games. Beliefs have been used to study best responses, belief learning, repeated-game reasoning, and higher-order beliefs (Nyarko and Schotter, 2002; Costa-Gomes and Weizsäcker, 2008; Danz et al., 2012; Aoyagi et al., 2024). Huck and Weizsäcker (2002) study beliefs about others' one-shot choice of lotteries, Hyndman et al. (2012) beliefs about others' play in repeated games, and Agranov and Detkova (2025) beliefs about others' belief updating; the use of an approach akin to an observer method is a feature shared by our belief experiment. Palfrey and Wang (2009) discuss belief elicitation in strategic settings more broadly and Friedman and Ward (2022) and Esteban-Casanelles and Gonçalves (2026) examine how beliefs about others' choices are affected by others' payoffs. Our experiment differs by focusing on the effect of different on- and off-path histories, where we elicit beliefs about the same continuation choice. This is what allows us to measure whether off-path moves are treated as informative about subsequent mistakes.

The paper also connects to work on stochastic choice, noisy beliefs, and cognitive frictions. Random utility, quantal response, perturbed utility, and costly information acquisition all imply that mistakes can be payoff-dependent (Luce, 1959; Block and Marschak, 1960; McKelvey and Palfrey, 1995, 1998; Mattsson and Weibull, 2002; Caplin and Dean, 2015; Matějka and McKay, 2015; Fudenberg et al., 2015; Strzalecki, 2023). An emerging literature on cognitive noise connects payoff-dependent mistakes to cognitive frictions (Enke and Graeber, 2023; Shubatt and Yang, 2026; Enke et al., 2026).

Experimental evidence also links smaller payoff differences to noisier choices and longer response times (Mosteller and Nogee, 1951; Krajbich et al., 2012; Clithero, 2018; Alós-Ferrer et al., 2021; Fudenberg et al., 2018; Spiliopoulos and Ortmann, 2018). Our one-shot and off-path results show that both choices and beliefs are payoff-sensitive even in very simple monetary problems. Our belief results also provide novel evidence that participants expect others’ choices to follow these patterns: larger payoff differences are expected to reduce mistakes, and larger off-path deviations are expected to increase subsequent mistakes. By comparing choice and belief patterns, we show that participants often exaggerate payoff sensitivity when predicting mistakes, yet underuse the selection information contained in off-path moves.

Finally, the paper relates to belief updating and inference from selected samples. The underinference we document is connected to classic evidence on belief-updating biases (Grether, 1980, 1992; Benjamin, 2019) and to recent work on cognitive constraints in belief updating (Agranov and Reshidi, 2026; Ba et al., 2025; Gonçalves et al., 2025). It is also related to selection neglect (Jehiel, 2018; Esponda and Vespa, 2018; Barron et al., 2024). In our setting, the selected sample is especially transparent: it is the set of participants who have just made an off-path move. We show that people recognise this selection, but underweight it. For perfection believers, the problem is even sharper, as they do not anticipate that an off-path move can occur. Their behaviour is related to belief revision after unforeseen events (Karni and Vierø, 2013; Becker et al., 2026), except that off-path moves in our experiment are objectively specified and only subjectively ruled out.

— —

The remainder of the paper proceeds as follows. **Section 2** presents the experimental design. **Section 3** discusses results regarding one-shot problems whereas **Section 4** focuses on first-stage choices of sequential problems, that is, off-path moves. **Section 5** reports on choices and beliefs on and off path and **Section 6** examines selection effects and selection neglect as possible mechanisms, whereas **Section 7** discusses heterogeneity regarding the beliefs about the likelihood of off-path moves. **Section 8** documents patterns regarding response times and **Section 9** concludes.

## 2. Experimental Design

We introduce an experimental paradigm that separately identifies choices, beliefs about choices, and beliefs following on- and off-path moves.

### 2.1. Overview of the Design

We elicit choices and beliefs across two experiments. The first experiment generates choice data in a simple monetary decision environment. The second experiment elicits beliefs, from a separate sample of participants, about the choices made in the first experiment.

In the choice experiment, participants first face one-shot binary choice problems. A problem is



denoted by  $P = \{h, l\}$ , where  $h > l$  are monetary amounts. We refer to  $h$  and  $l$  as the **maximum** and **minimum payoffs** of  $P$ , respectively. Choosing  $l$  is therefore a **suboptimal choice**.

Participants then face two-stage sequential choice problems. In the first stage, they choose between two binary choice problems,  $P = \{h, l\}$  and  $P' = \{h', l'\}$ . In the second stage, they choose one of the two alternatives in the problem selected in the first stage. Thus, if the participant selects  $P$ , the second-stage choice is between  $h$  and  $l$ ; if the participant selects  $P'$ , the second-stage choice is between  $h'$  and  $l'$ .

The sequential problems allow us to define on- and off-path moves in a transparent way. A first-stage choice is an **on-path** move if it selects the continuation problem with the larger maximum payoff and an **off-path** move otherwise. Thus, if  $h > h'$ , choosing  $P$  is on path and choosing  $P'$  is off path; if  $h < h'$ , the reverse is true. We restrict the on- versus off-path analysis to pairs with distinct maximum payoffs. Conditional on the first-stage choice, the second-stage choice is again a binary monetary choice, so choosing the lower payoff in the selected problem is suboptimal.

The belief experiment mirrors the choice experiment. Participants in the belief experiment are shown the same decision environments and report incentivised beliefs about the choices made by participants in the choice experiment. In the first one-shot part, they report beliefs about choices in each binary problem  $P$ . In the second sequential part, they first report beliefs about the first-stage choice between  $P$  and  $P'$ . They then report beliefs about the second-stage choice in each continuation problem, conditional on that problem having been selected.

The design has three key features. First, because all alternatives are own monetary payoffs, optimal and suboptimal choices are directly observable. In many dynamic-game experiments—as in centipede and ultimatum games—deviations from equilibrium can be instead seen as reflecting social preferences or fairness concerns (e.g., [Fehr and Schmidt, 1999](#); [McKelvey and Palfrey, 1992](#); [Andreoni and Blanchard, 2006](#); [Healy, 2024](#)). Whilst these are important considerations, they confound identification of suboptimal behaviour and, consequently, of beliefs about others following on- and off-path choices. In particular, by eliciting beliefs about other participants' one-shot and sequential decisions, the design deliberately avoids these confounds by shutting off this channel, whilst nevertheless casting light on beliefs about choices off path. Second, because the same binary problem can appear as a one-shot problem, as an on-path continuation problem, or as an off-path continuation problem, we can compare behaviour and beliefs at the same decision problem across path histories. That is, we can study how choices in  $P$  and beliefs about those choices depend on whether  $P$  is a one-shot problem, follows an off-path move, or an on-path move, when  $P$  was chosen out of  $\{P, P'\}$  with  $h < h'$  or  $h > h'$ . Third, by varying  $h$ ,  $l$ ,  $h'$ , and  $l'$ , we vary the payoff costs of suboptimal choices at all stages of the sequential problem. The **payoff spread** between alternatives ( $h - l$ ) captures the payoff cost of suboptimal second-stage or one-shot choices. The **difference in the maximum continuation payoffs** ( $h - h'$ ) is the payoff cost of first-stage off-path moves. We can also calculate the **difference in the minimum continuation payoffs** ( $l - l'$ ), which measures the

payoff cost if a participant were to make suboptimal second-stage choices. This allows us to study not only whether off-path histories affect continuation beliefs, but also whether these effects vary with the severity of the deviation.

When we refer to **expected continuation payoffs** in the analysis, we do not impose sequential rationality. Instead, we construct expected payoffs using observed average one-shot choice frequencies in the choice experiment, or reported one-shot average beliefs in the belief experiment. This distinction matters because the design is intended to measure how participants perceive mistakes, not to assume them away.

## 2.2. Implementation Details

We implemented the design using a simple interface in which monetary alternatives were displayed as balls inside boxes. Each ball corresponded to a monetary payment, denominated in pounds. This presentation was chosen to make the choice environment transparent and to minimise ambiguity about the payoff consequences of each action.

**Choice Experiment.** In the choice experiment, participants chose directly among monetary payoffs for themselves. In part 1, each binary choice problem  $P = \{h, l\}$  was displayed as a box containing two balls. Each ball displayed the corresponding bonus payment. [Figure 1](#) shows an example in which the participant chooses between  $l = £0.95$ , the left orange ball, and  $h = £1.05$ , the right teal ball.

We constructed nine binary problems by crossing three payoff spreads,  $h - l \in \{£0.10, £0.60, £1.20\}$ , with three payoff levels,  $l \in \{£0.45, £0.95, £1.35\}$ . In the example in [Figure 1](#), the payoff level  $l$  is £0.95 and the spread  $h - l$  is £0.10. The design systematically varies both the absolute payoff difference between the two alternatives and the level of the lower payoff. These dimensions allow us to test whether mistakes are more likely when payoff differences are small, as in models of noisy cognition and stochastic choice.

Part 2 used the same nine binary problems to construct two-stage sequential problems  $\{P, P'\}$ . In each round, a problem was generated by randomly combining two of the nine binary choice problems,  $P$  and  $P'$ . In the first stage, participants choose one of the two boxes corresponding to one of the two binary problems,  $P$  or  $P'$ . In the second stage, they choose a ball only within the box that they chose in the first stage, so between  $h$  or  $l$  if they chose  $P$  and between  $h'$  or  $l'$  if they chose  $P'$ . The second-stage choice determined the bonus payment for that round. [Figure 2](#) shows an example in which a participant in the first stage chose  $P = \text{'box 1'}$ . In the second stage, the participant chooses between the left yellow ball, paying  $l = £1.35$ , and the right ball, paying  $h = £1.45$ . If that participant had chosen  $P' = \text{'box 2'}$ , they would have faced a box with the payoffs from the one-shot example in [Figure 1](#).

The two-stage procedure allows for a clear interpretation of on- and off-path choices: choosing the box with the larger maximum payoff keeps the highest possible round payoff available, whereas




Figure 1: Choice Elicitation: One-Shot Choice

*Notes:* This figure provides a screenshot of the interface in a one-shot binary choice problem as shown to the participants in part 1 of the choice-elicitation experiment. Each ball depicted the corresponding bonus of money in pounds.



Figure 2: Choice Elicitation: Sequential Choice

*Notes:* This figure provides a screenshot of the interface in a two-stage binary choice problem as shown to the participants in part 2 of the choice-elicitation experiment. In the first stage, illustrated in panel (a), participants first chose a box, and in the second stage, illustrated in panel (b), participants chose a ball in the chosen box. The choice in the first stage is fixed and shown in the second stage. Each ball depicted the corresponding bonus of money in pounds.

choosing the other box rules it out. When discussing the effects of on- and off-path, we consider the pairs of binary choice problems ('boxes') with different maximum payoffs, that is,  $P = \{h, l\}, P' = \{h', l'\}$  such that  $\max P \neq \max P'$ .

Participants completed 36 incentivised rounds. There were 9 one-shot rounds in part 1 and 27 sequential rounds in part 2. In part 1, each of the nine binary problems appeared once, in random order. In part 2, participants faced a randomly selected set of pairs of binary problems. We randomised the left-right position of balls within boxes, the position of boxes on the screen, and the colour assigned to each ball. At the end of the experiment, one round was selected uniformly at random for payment. The participant received the bonus associated with the ball chosen in that round.

Box

£1.65

£0.45

What do you think is the probability that the other participant chose:

●  %
●  %

Next

Other Participants' Instructions

Your Instructions

Figure 3: Belief Elicitation: One-Shot Choice

*Notes:* This figure provides a screenshot of the interface in a one-shot binary choice problem as shown to the participants in part 1 of the belief-elicitation experiment. Each ball depicted the corresponding bonus of money in pounds.

Box 1

£2.15

£0.95

Box 2

£1.55

£0.95

What do you think is the probability that the other participant chose:

Box 1

 %

Box 2

 %

Next

Other Participants' Instructions

Your Instructions

You predicted the following probabilities of the other participant choosing Box 1:

5.0 %

Suppose you've been matched with one of the other participants who chose Box 1.

What do you think is the probability that the other participant chose:

●  %
●  %

You predicted the following probabilities of the other participant choosing Box 2:

95.0 %

Suppose you've been matched with one of the other participants who chose Box 2.

What do you think is the probability that the other participant chose:

●  %
●  %

Next

Other Participants' Instructions

Your Instructions

Figure 4: Belief Elicitation: Sequential Choice

*Notes:* This figure provides a screenshot of the interface in a two-stage binary choice problem as shown to the participants in part 2 of the belief-elicitation experiment. In the first stage, illustrated in panel (a), participants first reported beliefs about the chosen box, and in the second stage, illustrated in panel (b), participants reported beliefs about the chosen ball in each box, conditional on it being chosen. The belief reported in the first stage is fixed and shown in the second stage. Each ball depicted the corresponding bonus of money in pounds.

This experiment allows us to study how often participants make suboptimal choices and off-path moves, how these are related, and how they vary with payoff differences.

**Belief Experiment.** The belief experiment elicited beliefs from a separate sample of participants about the choices made in the choice experiment. Using a separate sample also avoids changing

choice behaviour through belief elicitation and makes belief participants outside observers of the choice environment (Hyndman et al., 2012; Palfrey and Wang, 2009). Participants first read the instructions of the choice experiment and were familiarised with the choice environment by making a practice choice. They were then told that their task was to predict the choices of participants who had previously completed that experiment and had to pass comprehension checks on the task.

Beliefs were incentivised using the binarised quadratic scoring rule (Hossain and Okui, 2013), determining the probability of receiving a bonus of £3.00. The instructions explained the payment rule in non-mathematical language, following Wilson and Vespa (2016). In line with recommendations by Danz et al. (2022), Healy and Kagel (2025), and Healy and Leo (2025), we indicated in the instructions that the “payment rule is designed so that you maximise your probability of winning the bonus payment if you tell us what you believe the probability is that a given ball or box was selected.”<sup>1</sup>

The belief experiment followed the same structure as the choice experiment. In part 1, participants reported beliefs about choices in the same nine one-shot binary problems. Figure 3 shows the interface. Participants could enter either probability in a pair of complementary text fields, with the other field updated automatically whenever the first is changed.

In part 2, participants reported beliefs about the sequential problems in two stages. First, they reported the probability that the choice-experiment participant selected each box, as in Figure 3a. Second, they reported beliefs about the choice of each ball in each box, conditional on that box having been selected, as per Figure 3b. That is, for each sequential problem, we elicited beliefs about both the first-stage path and the second-stage continuation choice following each possible path. The second-stage elicitation is analogous to a strategy-method elicitation of beliefs: participants reported their beliefs about choices *conditional* on each first-stage choice, including choices they considered unlikely. The interface reminded participants of their own first-stage belief before they reported the corresponding conditional beliefs. Figure 4 shows the corresponding interface.

## 2.3. Procedures

**Participants.** Both experiments were conducted on Prolific in June and October 2025. We recruited 450 participants in total, with 225 participants in each experiment. Eligibility was restricted to adult Prolific users located in the United Kingdom, fluent in English, and with an approval rate of at least 95%. Both experiments included comprehension checks that participants had to pass before proceeding to the paid rounds, as well as a text captcha. Participants could not take part in both experiments.

**Payment.** As is standard in experiments with multiple incentivised decisions, one round was selected uniformly at random for payment (Charness et al., 2021a; Azrieli et al., 2018). In the choice experiment, participants received a fixed payment of £2.55, plus the bonus based on their choices in

---

<sup>1</sup>For overviews of belief-elicitation methods in experiments, see Schotter and Trevino (2014), Schlag et al. (2015), Charness et al. (2021b), and Healy and Leo (2025).

one randomly selected round. Average total payment was £4.35, and the median completion time was 7 minutes. In the belief experiment, participants received a fixed payment of £3.00, plus potentially a bonus of £3.00, where the probability of receiving such bonus was based on the accuracy of one randomly selected belief report. Average total payment was £5.70, and the median completion time was 17 minutes.

**Randomisation.** For each participant, the order of the one-shot problems in part 1 and the order of the sequential problems in part 2 were uniformly randomised. Within each round, we randomised the position of the balls, the position of the boxes, and the colour assigned to each ball. Ball colours were drawn without replacement from a fixed palette within each display.

**Preregistration, Design, and Analysis.** The experimental design and the main hypotheses for the belief experiment were preregistered on AsPredicted (#232,256). The preregistration focused on beliefs because initial funding constraints limited the planned choice-experiment sample. As anticipated in the preregistration, we were subsequently able to increase the choice-experiment sample from 110 to 225 participants and the belief-experiment sample from 220 to 225 participants.

In the analysis below, we distinguish preregistered hypotheses, secondary analyses based on the design’s payoff variation, and exploratory analyses motivated by patterns observed in the data. The main preregistered hypotheses concern differences in beliefs across one-shot, on-path, and off-path environments, and how these differences vary with the severity of the first-stage deviation. The secondary analyses use the randomised payoff variation to study how choices and beliefs respond to payoff differences, consistent with models of stochastic choice and payoff-dependent mistakes.

### 3. Choice and Beliefs in One-Shot Settings

We first study choices and beliefs in one-shot settings. This section serves two purposes. First, it documents the baseline frequency of suboptimal choices in the simplest environment in the experiment. Second, it examines whether and how participants’ choices and beliefs about others’ choices depend on the payoff features of the choice problem.

#### 3.1. Hypotheses and Identification

Recall that one-shot problems are binary choices between two monetary amounts. Each problem is denoted by  $P = \{h, l\}$ , where  $h > l$ . The low-payoff alternative  $l$ —which we call the **payoff level**—takes one of three values, £0.45, £0.95, or £1.35. The high-payoff alternative is obtained by adding one of three **payoff spreads**  $h - l$ , £0.10, £0.60, or £1.20, to the payoff level. The first part of the choice experiment elicited choices in each of the nine resulting binary problems. The first part of the belief experiment elicited beliefs about those choices. In both experiments, the nine problems were presented in random order. Because payoff levels and payoff spreads are crossed, the design varies the two features independently.

It was not clear *ex ante* whether these payoff features would affect choices or beliefs. One possibility is that mistakes are payoff-dependent. Models reflecting cognitive costs or limitations, such as random utility (Luce, 1959; Block and Marschak, 1960; McFadden, 1974), additive perturbed utility (Mattsson and Weibull, 2002; Fudenberg et al., 2015), or costly information acquisition (Caplin and Dean, 2015; Matějka and McKay, 2015) all imply that choices can become noisier when the payoff cost of a mistake is smaller.<sup>2</sup> This prediction is also consistent with experimental evidence that choices are noisier when payoff differences are smaller (e.g., Mosteller and Nogee, 1951; Clithero, 2018), even in strategic contexts (McKelvey and Palfrey, 1995; Friedman and Ward, 2022; Almond and Martin, 2024). Under this view, a larger payoff spread should reduce the frequency of mistakes. A larger payoff level may instead increase the frequency of suboptimal choices because, holding the spread fixed, it reduces the relative payoff difference between the two alternatives.<sup>3</sup> More generally, if participants believe that others make payoff-sensitive mistakes, reported beliefs should vary with the same payoff features even when their own choices are nearly deterministic.

A second possibility is that payoff variation has little effect in our setting. The monetary amounts are displayed directly, and participants do not need to assess, compute, or learn the value of the alternatives. This distinguishes our setting from many environments in which payoff-dependent mistakes arise due to more significant cognitive frictions, where payoffs to alternatives must be evaluated or inferred. The simplicity of the one-shot problems therefore makes a null effect plausible. This *ex-ante* ambiguity was also reflected in our own expectations before observing the data. We test the following hypothesis:

**Hypothesis 1.** In one-shot choice problems, the frequency of suboptimal choices and the average belief about suboptimal choices (a) decrease in the payoff spread and (b) increase in the payoff level.

### 3.2. Suboptimal Choices and Beliefs about Suboptimal Choices

Suboptimal choices in one-shot problems are rare and are also rarely expected. Across all one-shot problems, participants chose the lower payoff in 2.37% of choices. Participants in the belief experiment expected substantially more mistakes. The average reported belief that another participant chose the lower payoff was 14.17%.

Both choices and beliefs vary with the payoff spread. Figure 5 shows the frequency of suboptimal choices in blue circles and the average reported belief about suboptimal choices in orange diamonds. Panel (a) plots these outcomes against the payoff spread, whereas panel (b) does so against the payoff level. Both choices and beliefs decrease with the payoff spread, but neither choices nor beliefs vary meaningfully with the payoff level. The corresponding regressions, which condition on both payoff

<sup>2</sup>For broader discussions of attention and costly information acquisition see Caplin (2016) and Caplin et al. (2022). See Strzalecki (2023) for a comprehensive textbook treatment of stochastic choice models.

<sup>3</sup>This force is absent in models in which choice probabilities depend only on absolute payoff differences. This is the case, for example, in additive perturbed-utility models, random-utility models with additively separable payoff shocks, and costly-information-acquisition models with additively separable information costs. A level effect can arise instead when choice noise scales with payoff levels, when participants evaluate payoff differences in relative terms, or when costs to information are multiplicative rather than additive.






Figure 5: Payoffs, Choices, and Beliefs in One-Shot Binary Choice Problems

*Notes:* This figure shows the frequency of suboptimal choices (blue circles) and the average reported belief about suboptimal choices (orange diamonds) in one-shot binary choice problems, both expressed in percentage points. Choices are from part 1 of the choice-elicitation experiment and beliefs from part 1 of the belief-elicitation experiment. Panels (a) and (b), respectively, show how choices and beliefs vary with the payoff spread and the payoff level. The payoff spread is the difference between the two monetary amounts in the choice problem. The payoff level is the minimum payoff in the choice problem. Payoffs are measured in pounds. Whiskers correspond to 95% confidence intervals using standard errors clustered at the participant level. Lines correspond to the best linear fit from a univariate regression.

features simultaneously, are reported in Table 4 in Appendix A.1. We therefore find support for Hypothesis 1a, but not for Hypothesis 1b.

Two additional patterns are important for the rest of the paper. First, participants substantially overestimate the frequency of mistakes. This overestimation is present even though the choice problem is transparent and mistakes are empirically infrequent. This can be seen as expressing a tendency to “not act as though the hypothesis that other people are utility maximizers is absolutely dependable” (Rosenthal, 1981), which has been experimentally documented in Beard and Beil (1994). Second, beliefs are more sensitive to payoff spreads than choices are. This pattern foreshadows the central theme of the paper: participants recognise that mistakes are more likely when they are less costly, but their beliefs exaggerate the payoff sensitivity of those mistakes. This can be seen in Figure 5, where the relationship between beliefs and payoff spreads is steeper than the corresponding relationship for choices. The regression results in Table 4 confirm the same pattern. The difference is driven by beliefs about low-spread problems. When the payoff spread is small, participants substantially overestimate the frequency of suboptimal choices. As the payoff spread increases, the gap between beliefs and choice frequencies narrows and, as a consequence, the belief bias, the absolute difference between reported beliefs and choice frequencies, decreases with the payoff spread.



## 4. The Off-Path

We now study first-stage choices in the sequential problems. In particular, we examine whether participants choose the on-path or off-path continuation problem, and whether choices and beliefs about these choices respond to the payoff cost of moving off path.

### 4.1. Hypotheses and Identification

The second part of the experiments allows us to study choices to move off path and beliefs about such choices. Participants faced 27 sequential choice problems, each constructed as a pair of binary choice problems. In the first stage, participants chose one of the two binary problems. In the second stage, they chose one of the two alternatives in the selected problem. In the belief experiment, participants reported beliefs about both the first-stage choice and the subsequent choice within each continuation problem.

We use the randomised variation in the payoffs of the sequential binary choice problems to study whether off-path moves are payoff-dependent. For each sequential problem, the off-path move is the first-stage choice that selects the continuation problem with the lower maximum payoff. The payoff cost of moving off path can therefore be measured by the difference between the maximum continuation payoff on path and the maximum continuation payoff off path. This measure corresponds to the foregone payoff under sequentially rational continuation choices. We also consider the difference in expected continuation payoffs, which allows for the possibility that participants expect to make mistakes in the continuation problem. To construct expected continuation payoffs, we use average one-shot choice frequencies in the choice experiment and average one-shot beliefs in the belief experiment.

The effect of payoff differences on off-path moves was ambiguous *ex ante*. On the one hand, the design is transparent. All payoffs are shown directly as monetary amounts, so participants may recognise immediately which continuation problem preserves the highest possible payoff. In other words, the simplicity of this ‘what-you-see-is-what-you-get’ design, framed in monetary amounts, was meant to make it clear whether a particular continuation binary choice problem was optimal or not, that is, on or off path. This makes a null effect plausible.

On the other hand, the same stochastic-choice and costly-cognition models discussed in [Section 3](#) predict that the probability of a suboptimal choice should fall when the foregone payoff associated with that choice increases ([McKelvey and Palfrey, 1998](#); [Fudenberg et al., 2015](#); [Strzalecki, 2023](#)). Under this view, off-path moves should be less frequent when the difference in continuation values is larger. The argument applies when one considers the continuation value as being given by the maximum continuation payoff—corresponding to sequentially rational choices—or the expected continuation payoffs, allowing for noise in choices. Moreover, if participants’ beliefs move in line with payoff-dependent mistakes, beliefs about off-path moves should respond in the same direction.

In short, since our design pushes the boundaries of the applicability of standard insights from cognitive cost models, it was ex-ante unclear whether payoffs would influence choices and beliefs about off-path moves in this simple setting. Our analysis tests exactly this question, summarised in the following hypothesis:

**Hypothesis 2.** In sequential choice problems, the frequency of off-path moves and the average belief about off-path moves decrease in the difference in continuation values.

## 4.2. Choosing and Believing in the Off-Path


We first document that participants' behaviour and their perception of what is on path and off path coincide with optimal sequential decision-making: off-path moves are uncommon, and participants believe them to be uncommon. Participants in the choice experiment moved off path in 5.73% of first-stage choices. Participants in the belief experiment assigned an average probability of 16.96% to the off-path box being chosen. Therefore, as in the one-shot problems, beliefs overstate the frequency of mistakes. Nevertheless, both choices and beliefs classify off-path moves as unlikely. Moreover, in 36% of first-stage belief reports, participants assigned probability zero to the off-path box being chosen. Similarly, in a majority of first-stage belief reports, participants assigned probability below 10% to the off-path box being chosen.

Off-path choices and beliefs are payoff-dependent. **Figure 6a** shows the frequency of off-path moves and the average belief about off-path moves as a function of the difference in maximum continuation payoffs. Both series decline as the difference in maximum payoffs increases across the fourteen differences in maximum payoffs in our design. This pattern indicates that participants are less likely to move off path, and expect others to be less likely to move off path, when the off-path move entails a larger relative decrease in payoff.

**Figure 6b** repeats the comparison using the difference in expected continuation payoffs. When considering expected continuation payoffs instead, we incorporate the possibility and anticipation of mistakes in the continuation binary choice problem. The results are remarkably similar. In short, the evidence supports **Hypothesis 2**, whether continuation values are measured by maximum payoffs or by expected payoffs. The corresponding regressions are reported in **Table 5** in **Appendix A.2**.

The data also reveal an unanticipated role for minimum payoffs: off-path moves are more frequent when the minimum payoff in the off-path continuation problem is high relative to the minimum payoff in the on-path continuation problem. **Figure 6c** shows this pattern for the seven differences in minimum payoffs. When the difference in minimum continuation payoffs is negative, the worst-case payoff is higher off path than on path. Participants are then more likely to make an off-path move and participants in the belief experiment anticipate this pattern, as beliefs about off-path moves also decrease with the difference in minimum payoffs.

The negative relationship between the difference in minimum payoffs and off-path choices is not



(a) Difference in Maximum Payoffs (b) Difference in Expected Payoffs (c) Difference in Minimum Payoffs

Figure 6: Payoffs, Choices, and Beliefs about Off-Path Moves

*Notes:* This figure shows the frequency of off-path moves (blue circles) and the average reported belief about off-path moves (orange diamonds) in the first stage of sequential binary choice problems, both expressed in percentage points. Choices are from part 2 of the choice-elicitation experiment and beliefs from part 2 of the belief-elicitation experiment. Panels (a), (b), and (c), respectively, show how choices and beliefs vary with the difference in maximum, expected, and minimum continuation payoffs between the on-path and off-path binary choice problems. Payoffs are measured in pounds. Whiskers correspond to 95% confidence intervals using standard errors clustered at the participant level. Lines correspond to the best linear fit from a univariate regression.

explained by the correlation between minimum and expected continuation payoffs.<sup>4</sup> Indeed, this relationship remains when controlling for the difference in expected continuation payoffs. The corresponding regressions are reported in Table 6 in Appendix A.2.

One interpretation is that some participants treat the first-stage decision as a choice over exposure to future choice mistakes. When the off-path problem has a higher minimum payoff, choosing it sacrifices a potential higher maximum payoff, but reduces the payoff loss from a possible future mistake. In this sense, the first-stage decision resembles a choice over continuation menus. This result is reminiscent of models of commitment and self-control, where agents may prefer to restrict future choice sets because flexibility exposes them to temptation or self-control costs (Strotz, 1955; Gul and Pesendorfer, 2001; Toussaert, 2018). The analogy is not literal: in our setting, the low-payoff alternative is dominated rather than tempting. The relevant concern is instead that the participant may later make a dominated choice, or may need to exert attention to avoid doing so. The result is therefore closer to commitment against future mistakes, or to noisy foresight (Chakraborty and Kendall, 2026), than to self-control in the standard temptation sense.

This interpretation also clarifies the connection to uncertainty aversion. The relevant uncertainty is not uncertainty about an external state, as in standard ambiguity aversion (Gilboa and Schmeidler,

<sup>4</sup>When the difference in minimum continuation payoffs is negative, the payoff spread is mechanically larger in the on-path problem than in the off-path problem. Since larger spreads are associated with fewer suboptimal choices in one-shot problems, this could mechanically increase the expected continuation payoff advantage of the on-path problem.

1989), but uncertainty about future choice behaviour. In this sense, the result is closer to aversion to unpredictability in dynamic choice (Ke, 2019). The magnitudes suggest that this behaviour cannot be rationalised by the empirical rate of one-shot mistakes alone.<sup>5</sup> We therefore treat the minimum-payoff result as evidence that some first-stage choices respond not only to expected continuation values, but also to the distribution of possible continuation payoffs. The fact that belief participants predict the same minimum-payoff pattern suggests that they also expect others to treat first-stage choices partly as choices over exposure to future mistakes.

## 5. Choices and Beliefs On- versus Off-Path

This section studies our main object of interest: how choices and beliefs in a continuation problem change when that same problem follows an on-path or off-path move. We compare three environments. A binary problem can appear as a one-shot problem, as a continuation problem following an on-path move, or as a continuation problem following an off-path move. Our outcomes are the frequency of suboptimal continuation choices and beliefs about those choices.

### 5.1. Hypotheses and Identification

The random pairing of binary problems in the sequential part of the experiments implies that the same continuation problem can appear on path or off path. Our identification uses this variation to compare choices and beliefs for the same binary problem across path histories. In particular, we compare behaviour and beliefs when a binary problem appears as a one-shot problem, when it follows an on-path move, and when it follows an off-path move. This comparison identifies whether the path by which a problem is reached affects subsequent choices and beliefs about those choices.

A classical subgame-perfect argument predicts no such path effect. Sequential rationality requires optimal continuation behaviour after every history, including histories that are not reached in equilibrium (Selten, 1975; Kreps and Wilson, 1982). The textbook justification is that an off-path deviation can be treated “as the result of an extremely unlikely ‘mistake’ that is unlikely to occur again” (Mas-Colell et al., 1995, p. 282). Under this interpretation, an off-path move is a tremble rather than evidence of a persistently less accurate decision-maker. Choices should therefore be no less optimal off path than on path or in one-shot problems. Participants who reason in this way should also hold similar beliefs about continuation choices on and off path.

A different prediction arises if participants differ in their propensity to make mistakes. For example, suppose some participants usually choose optimally, whereas others are more inattentive or more prone to payoff-dependent mistakes. Conditioning on an off-path move then creates a selection effect: after an off-path move, the continuation problem is more likely to be faced by a participant

---

<sup>5</sup>Consider cases in which the minimum payoff following an off-path move is strictly higher than the minimum payoff following an on-path move. For any such sequential choice problem, the continuation mistake probability that would make a participant weakly prefer the off-path problem is strictly greater than 16.67%. By contrast, in the data, the frequency of suboptimal choices in any binary one-shot problem is never higher than 4.44%.

with a higher mistake propensity.<sup>6</sup> Participants in the belief experiment may anticipate this selection and assign a higher probability to suboptimal continuation choices off path, compared to choices following a move on path.

This selection logic also generates comparative statics. If mistakes are payoff-dependent, then a larger first-stage deviation should be more diagnostic of the participant’s mistake propensity. Such a prediction requires both cognitive costs or frictions leading to payoff-sensitive mistakes, as in quantal-response or stochastic-choice models (McKelvey and Palfrey, 1998; Fudenberg et al., 2015), and heterogeneity in the tendency to make those mistakes. Under this view, off-path histories should increase suboptimal continuation choices, and the effect should be larger when the off-path move forgoes a larger continuation value. Beliefs should move in the same direction if participants account for this selection.

There is also a possible force in the opposite direction. If recognising that one has made a mistake makes a participant more alert, then suboptimal choices may become less frequent after an off-path move.<sup>7</sup> The sign of the path effect is therefore an empirical question.

Our design is uniquely well suited to distinguish these possibilities. In many dynamic-game experiments, deviations from subgame perfection can reflect social preferences, fairness concerns, or beliefs about others’ preferences rather than mistakes alone (e.g., Fehr and Schmidt, 1999; McKelvey and Palfrey, 1992; Andreoni and Blanchard, 2006; Healy, 2024). Our setting removes this particular confound because all choices are between own monetary payoffs. Moreover, the random payoff variation allows us to test whether the informativeness of on- and off-path histories varies with the foregone payoffs associated with first-stage deviations. Our guiding hypothesis is the following:

**Hypothesis 3.** In continuation problems, the frequency of suboptimal choices and the average belief about suboptimal choices (a) are higher following an off-path move and lower following an on-path move, and (b) these effects increase with the difference in continuation values.

## 5.2. Path Effects on Choices and Beliefs

We now study the frequency of suboptimal choices and the average reported belief about suboptimal choices in three environments: one-shot problems, continuation problems following on-path moves, and continuation problems following off-path moves.

We quantify these effects in Table 1. Columns (1) and (4) estimate the regression

$$y_i = \beta_0 + \beta_1 [\text{Off Path}]_i + \beta_2 [\text{On Path}]_i + \epsilon_i, \quad (1)$$

where  $y_i$  is either an indicator for a suboptimal choice, coded as 0 or 100, or the reported belief

<sup>6</sup>Similar selection effects appear in models with level- $k$  players (Nagel, 1995; Stahl and Wilson, 1995; Camerer et al., 2004) in extensive-form games (Ho and Su, 2013; Lin and Palfrey, 2024; Fong et al., 2025). Such models, however, would not be able to rationalise payoff-dependent mistakes in our data.

<sup>7</sup>We did not view this force as the most likely one ex ante, but it was suggested in several conversations before we analysed the data.

|                            | Chose Suboptimally   |                      |                      |                      |                      |                      |
|----------------------------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
|                            | Choice               |                      |                      | Belief               |                      |                      |
|                            | (1)                  | (2)                  | (3)                  | (4)                  | (5)                  | (6)                  |
| Off Path                   | 10.314***<br>(2.920) | 0.208<br>(2.334)     | 0.286<br>(2.317)     | 4.342***<br>(0.457)  | 1.614***<br>(0.503)  | 2.091***<br>(0.504)  |
| On Path                    | -0.594<br>(0.449)    | -0.257<br>(0.544)    | -0.277<br>(0.542)    | -1.367***<br>(0.422) | 0.158<br>(0.488)     | -0.090<br>(0.486)    |
| Off Path x Diff Max Payoff | –                    | 22.554***<br>(6.646) | –                    | –                    | 3.590***<br>(0.440)  | –                    |
| On Path x Diff Max Payoff  | –                    | -0.434<br>(0.328)    | –                    | –                    | -2.006***<br>(0.409) | –                    |
| Off Path x Diff Exp Payoff | –                    | –                    | 22.676***<br>(6.701) | –                    | –                    | 3.131***<br>(0.441)  |
| On Path x Diff Exp Payoff  | –                    | –                    | -0.410<br>(0.327)    | –                    | –                    | -1.775***<br>(0.421) |
| Intercept                  | 2.370***<br>(0.637)  | 2.370***<br>(0.637)  | 2.370***<br>(0.637)  | 14.168***<br>(0.955) | 14.168***<br>(0.956) | 14.168***<br>(0.956) |
| R <sup>2</sup>             | 0.02                 | 0.04                 | 0.04                 | 0.02                 | 0.02                 | 0.02                 |
| N                          | 7937                 | 7937                 | 7937                 | 13835                | 13835                | 13835                |

Table 1: Choices and Beliefs Conditional on On- versus Off-Path Moves

*Notes:* This table shows how the frequency of suboptimal choices and beliefs about suboptimal choices, both expressed in percentage points, depend on whether the binary choice problem follows an on- or off-path move and on the difference in maximum or expected continuation payoffs. The data comprise one-shot binary choice problems in part 1 of the experiments and second-stage choices and beliefs from the sequential binary choice problems in part 2 of the experiments. Off Path and On Path are indicator variables equal to one if the binary choice problem follows an off- or on-path move, respectively, and zero otherwise. Diff Max/Exp Payoff denotes the difference in maximum/expected continuation payoffs across the on- and off-path binary choice problems. Payoffs are measured in pounds. Choice data are from the choice-elicitation experiment and belief data are from the belief-elicitation experiment. Standard errors clustered at the participant level are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.

that a suboptimal choice occurred, expressed in percentage points.<sup>8</sup> The variable  $[\text{Off Path}]_i$  equals one if the binary problem follows an off-path move in the sequential environment, and  $[\text{On Path}]_i$  equals one if it follows an on-path move. When  $[\text{Off Path}]_i = [\text{On Path}]_i = 0$ , the observation refers to the same binary problem presented as a one-shot problem. Thus,  $\beta_0$  captures the one-shot baseline,  $\beta_1$  captures the average off-path effect, and  $\beta_2$  captures the average on-path effect. The remaining columns interact the path indicators with the difference in maximum or expected continuation payoffs, providing a direct test of **Hypothesis 3b**.

Our main result concerns the frequency of suboptimal choices and the corresponding beliefs depending on the path by which a problem is reached. Columns (1) and (4) of **Table 1** support **Hypothesis 3a**. In the choice data, column (1) shows that the frequency of suboptimal choices increases by more than 10 percentage points when a binary problem follows an off-path move rather than appearing

<sup>8</sup>We code choices as 0 or 100 to make the choice coefficients directly comparable to belief coefficients, which are expressed in percentage terms.



as a one-shot problem. Choices following on-path moves are not significantly less likely to be suboptimal than one-shot choices. This may partially reflect a ceiling effect, given the low frequency of suboptimal choices in the one-shot setting. Therefore, it is off-path moves that reveal a strong selection effect in the choice data: participants who move off path in the first stage are substantially more likely to choose suboptimally in the continuation problem.

Column (4) shows that participants in the belief experiment also condition their beliefs on the path by which the continuation problem is reached. They assign a 4.34 percentage point higher probability to suboptimal choices off path and a 1.37 percentage point lower probability to suboptimal choices on path. This pattern is consistent with participants recognising that off-path moves select participants with higher mistake propensities. As in the one-shot setting, however, beliefs overstate the baseline level of suboptimal choices: the one-shot intercept is 2.37% in the choice data against 14.17% in the belief data (columns (1) and (4)).

Columns (2), (3), (5), and (6) of [Table 1](#) address [Hypothesis 3b](#). In the choice data, column (2) shows that the off-path effect is larger when the first-stage deviation entails a larger difference in maximum payoffs. Column (3) repeats the exercise using the difference in expected continuation payoffs instead of the difference in maximum continuation payoffs. The resulting estimates are not only qualitatively but also quantitatively very similar. In sum, more severe off-path deviations—as captured by larger differences in continuation payoffs—are followed by larger increases in suboptimal continuation choices.

Beliefs respond in the same direction, but less strongly. Columns (5) and (6) of [Table 1](#) show that participants expect more suboptimal continuation choices after larger off-path deviations. They also expect fewer suboptimal continuation choices after larger on-path payoff differences. However, the belief coefficients are substantially smaller than the corresponding choice coefficients for off-path histories. Participants therefore understand that off-path moves are informative about continuation behaviour, but underestimate how informative they are. This underinference is notable because beliefs are often more responsive than choices to payoff variation elsewhere in the experiment. Here, by contrast, beliefs move in the right direction but not enough.

## 6. Selection Effects and Selection Neglect

The previous section established that whilst participants in the choice-elicitation experiment make more suboptimal choices following an off-path move, beliefs about such choices seem to fail to adjust sufficiently strongly. Moves off path, while very unlikely, are also very informative. Participants who choose the suboptimal path are substantially more likely to make subsequent suboptimal choices. Belief participants partly recognise this informativeness, but do not fully appreciate it.

In this section, we study the mechanism behind this pattern. We first show that participants differ in their propensity to make suboptimal choices, and that this heterogeneity generates selection into off-path histories. We then quantify the extent to which beliefs account for this selection. The






Figure 7: Suboptimal Choices in One-Shot and Sequential Choice Problems

*Notes:* This figure shows the frequency—expressed in percentage points—of off-path moves (light-blue circles) and suboptimal choices in the selected continuation binary choice problems (dark-blue diamonds) in part 2 of the choice-elicitation experiment, conditional on the number of suboptimal choices the participant made in one-shot problems in part 1. Whiskers correspond to 95% confidence intervals using heteroscedasticity-robust standard errors.


evidence indicates selection neglect: beliefs move in the right direction, but substantially underinfer from the information contained in off-path moves.

## 6.1. Suboptimal Choices and Selection into the Off-Path

The selection interpretation requires heterogeneity in mistake propensities. If some participants are more likely than others to make suboptimal choices, then those participants should also be more likely to move off path in the first stage. Off-path histories would then select participants with higher continuation mistake rates.

Figure 7 provides direct evidence for this mechanism. It groups participants by the number of suboptimal choices they made in one-shot problems in part 1 of the choice experiment. Participants who made more one-shot mistakes are also more likely to move off path in part 2 and more likely to make suboptimal choices in the selected continuation problems. Moreover, we find that the participant-level frequency of one-shot mistakes is also strongly correlated with both the frequency of off-path moves and the frequency of suboptimal continuation choices: the corresponding pairwise correlations are 0.64 and 0.71, respectively, and both are statistically significant ( $p$ -value  $< .001$ ).

Participant-level heterogeneity in suboptimal choices is limited but significant. Figure 8a shows the cumulative distribution of participant-level average frequencies of suboptimal choices in one-shot binary choice problems, following on-path moves, and following off-path moves. Most participants always optimally choose the alternative with the higher payoff. Among those who make mistakes, however, mistake rates vary. A  $\chi^2$  test rejects the null that all participants have the same rate of suboptimal choices, whether we use only part 1 choices, only part 2 binary choices, or all binary choice problems ( $p$ -value  $< 0.001$  in all cases). The figure also shows selection into off-path histories.



(a) Distribution of Suboptimal Choice Frequency

(b) Distribution of Estimated Logit Parameter

Figure 8: Participant Heterogeneity in Choices

*Notes:* Panel (a) shows the cumulative distribution of participant-level frequencies of suboptimal choices in one-shot binary choice problems (solid blue line), following on-path moves (dotted light-blue line), and following off-path moves (dashed dark-blue line). Values are expressed in percentage points. One-shot choices are from part 1 of the choice-elicitation experiment. Choices following on- and off-path moves are from the second stage of sequential choice problems in part 2.

Panel (b) shows the cumulative distribution of participant-specific estimated logit parameters  $\lambda \in [0, 1]$ . When choosing between monetary amounts  $h > l$ , the model is given by  $\mathbb{P}(\text{Chose Suboptimally}) = (1 + \exp((h - l)(1 - \lambda)/\lambda))^{-1}$ . Parameters are estimated by maximum likelihood for each participant using all binary choice problems in the choice-elicitation experiment.

The distribution of participant-level mistake rates following off-path moves (dashed dark blue) first-order stochastically dominates the corresponding distributions for one-shot problems and on-path continuation problems. This is consistent with the correlations above: participants who are more likely to make binary-choice mistakes are also more likely to move off path.

Heterogeneity in cognitive costs or limitations is also crucial to explain why the path effect varies with the payoff cost of the first-stage deviation. Indeed, although heterogeneity in mistake propensities explains why off-path histories are followed by more suboptimal continuation choices, it does not, by itself, account for how the increase depends on the difference in continuation values. To capture payoff-dependent mistakes, we estimate a simple participant-level logit model. For a participant choosing between  $h$  and  $l$ , with  $h > l$ , we assume

$$\mathbb{P}(\text{Chose Suboptimally}) = \left( 1 + \exp \left( (h - l) \frac{1 - \lambda}{\lambda} \right) \right)^{-1},$$

where  $\lambda \in [0, 1]$ . The parameter  $\lambda$  indexes choice noise: as  $\lambda \rightarrow 0$ , choices become optimal; as  $\lambda \rightarrow 1$ , choices become uniformly random. This specification is a reduced-form way to capture payoff-sensitive mistakes, as in stochastic-choice and quantal-response models, often motivated as capturing cognitive frictions (McKelvey and Palfrey, 1995, 1998; Fudenberg et al., 2015; Strzalecki, 2023).

Figure 8b shows the cumulative distribution of participant-level maximum-likelihood estimates of  $\lambda$ . The distribution again reflects limited heterogeneity, because many participants always choose

optimally. Nevertheless, among participants who make mistakes, there is meaningful variation in estimated choice noise. Even considering only participants who made mistakes ( $\lambda \in (0, 1)$ ), a bootstrap likelihood-ratio test rejects a homogeneous logit model with a common  $\lambda$  against a model with participant-specific  $\lambda$  parameters ( $p$ -value  $< 0.001$ ).

Despite the limited individual-level data, this simple model predicts strong selection effects. Using fitted values, the predicted frequency of suboptimal choices is 3.03% in one-shot problems. It decreases to 1.47% conditional on an on-path move and increases to 36.00% conditional on an off-path move.<sup>9</sup> Therefore, the model captures the central mechanism: participants make payoff-dependent mistakes, and participants with higher estimated mistake propensities are more likely to select into off-path histories.

We conclude that the choice data support a selection interpretation. Off-path moves are informative because they are disproportionately made by participants who are more likely to make suboptimal choices in binary choice problems.

## 6.2. Belief Updating and Off-Path Selection Neglect

Whilst the choice data reveal substantial selection into off-path histories, the belief data show that participants only partially recognise this selection. We now seek to quantify the extent of this underinference.

From [Table 1](#), the data indicate that, following an off-path move, the empirical frequency of suboptimal continuation choices increases sharply relative to the one-shot benchmark, whereas beliefs are distinctly less affected. Note that this is not simply a general insensitivity of beliefs to the details of the problem. In earlier sections, beliefs often responded even more strongly to variations in payoffs than choices did. Here, by contrast, beliefs respond in the right direction but underreact to the information contained in the path.

To systematically investigate and quantify this pattern, we adapt a Grether-style belief-updating regression ([Grether, 1980, 1992](#)). The standard specification relates posterior log-odds of an event A relative to event B realising to prior log-odds and the log-likelihood ratio of an observed signal:

$$\ln \left( \frac{\text{Posterior}(\text{Event A} \mid \text{Signal})}{\text{Posterior}(\text{Event B} \mid \text{Signal})} \right) = \beta_0 \ln \left( \frac{\text{Prior}(\text{Event A})}{\text{Prior}(\text{Event B})} \right) + \beta_1 \text{Log-Likelihood}(\text{Signal}), \quad (2)$$

where the log-likelihood of the signal, i.e., its precision, is defined as usual:  $\text{Log-Likelihood}(\text{Signal}) := \ln \left( \frac{\mathbb{P}(\text{Signal} \mid \text{Event A})}{\mathbb{P}(\text{Signal} \mid \text{Event B})} \right)$ . Bayesian updating corresponds to  $\beta_0 = \beta_1 = 1$ . Base-rate neglect corresponds to  $\beta_0 < 1$ , and underinference from the signal corresponds to  $\beta_1 < 1$ . This regression is widely used to study belief-updating patterns; see [Benjamin \(2019\)](#) for a review. Recent work also provides formal foundations for updating rules of this form under monotone likelihood-ratio restrictions (cf. [Chan, 2026](#); [Cripps, 2025](#)).

<sup>9</sup>The predicted off-path selection effect is larger than the empirical effect. This is unsurprising, since off-path moves are infrequent and participant-level logit parameters are necessarily very imprecisely estimated.

In our setting, the signal is the whether the on- or off-path is, whereas the event of interest is whether the participant chooses suboptimally in the chosen continuation problem. Thus, within a particular sequential choice problem, where participants choose first between  $P, P'$  and then between the alternatives in the chosen problem, Event A is a suboptimal continuation choice, Event B is an optimal continuation choice, and the signal is whether the continuation problem follows an off-path or on-path move, that is, to whether  $P$  or  $P'$  was chosen in the first stage. For empirical choice frequencies, equation (2) holds mechanically with  $\beta_0 = \beta_1 = 1$  when priors, posteriors, and likelihood ratios are all computed from the same choice data. The question is whether reported beliefs respond to the empirical informativeness of the path in the same way.

Unlike in standard belief-updating experiments, priors and signal precisions are not induced. In this sense, beliefs are truly subjective, and participants are not told the empirical likelihood ratio of an on- or off-path move. We therefore use the choice data to estimate the empirical log-likelihood of each path move and use reported beliefs to construct the corresponding subjective prior. Specifically, for each sequential problem, we compute the empirical log-likelihood of the observed path move from the difference between posterior and prior log-odds in the choice data. We obtain Log-Likelihood(Signal) from the identity

$$\text{Log-Likelihood(Signal)} \equiv \ln \left( \frac{\mathbb{P}(\text{Chose Suboptimally} \mid \text{Signal})}{\mathbb{P}(\text{Chose Optimally} \mid \text{Signal})} \right) - \ln \left( \frac{\mathbb{P}(\text{Chose Suboptimally})}{\mathbb{P}(\text{Chose Optimally})} \right),$$

where the signal is whether the first-stage move in the sequential decision problem is on path or off path.<sup>10</sup> We then derive the subjective prior implied by each participant's reported path belief and conditional continuation beliefs:

$$\begin{aligned} \text{Prior(Chose Suboptimally)} := & \text{Belief(Off-Path Move)} \times \text{Posterior(Chose Suboptimally} \mid \text{Off-Path Move)} \\ & + \text{Belief(On-Path Move)} \times \text{Posterior(Chose Suboptimally} \mid \text{On-Path Move)}. \end{aligned} \quad (3)$$

Finally, we estimate equation (2) and variants that interact the log-likelihood with an off-path indicator [Off-Path] for whether the move is an off-path move ([Off-Path] = 1) or on-path ([Off-Path] = 0) and with the difference in maximum continuation payoffs across the on- and off-path binary choice problems.

**Table 2** reports the estimates. Participants substantially underinfer from path moves. Column (1) shows that the coefficient on Log-Likelihood(Path Move) is approximately 0.11, far below the Bayesian benchmark of one. Beliefs therefore respond to the empirical informativeness of the path, but only weakly. This degree of underinference is much larger than what is commonly found in standard belief-updating experiments using balls-and-urn paradigms (Benjamin, 2019). Column (2) shows that underinference is especially severe for off-path moves. The interaction between Log-Likelihood(Path Move) and Off Path is negative, indicating that beliefs respond less to the informativeness of off-path signals than to the informativeness of on-path signals. This pattern is consistent

<sup>10</sup>We winsorise frequencies by setting a minimum of 0.1% and a maximum of 99.9% so that log-odds are well defined. The results are qualitatively similar for other winsorisation thresholds, such as 0.25% and 0.5%.

|                                             | Posterior(Chose Suboptimally) |                      |                      |
|---------------------------------------------|-------------------------------|----------------------|----------------------|
|                                             | (1)                           | (2)                  | (3)                  |
| Prior(Chose Suboptimally)                   | 0.967***<br>(0.005)           | 0.956***<br>(0.006)  | 0.960***<br>(0.006)  |
| Log-Likelihood(Path Move)                   | 0.107***<br>(0.009)           | 0.365***<br>(0.034)  | 0.238***<br>(0.029)  |
| Log-Likelihood(Path Move) x Off Path        | –                             | -0.275***<br>(0.030) | -0.256***<br>(0.029) |
| Log-Likelihood(Path Move) x Diff Max Payoff | –                             | –                    | 0.123***<br>(0.016)  |
| N                                           | 11810                         | 11810                | 11810                |

Table 2: Belief Updating from Path Moves

*Notes:* This table reports estimates of the belief-updating regression in equation (2). Column (1) reports the baseline specification, and columns (2) and (3) add interaction terms. Posterior(Chose Suboptimally) is the reported belief about another participant’s second-stage choice in the sequential choice problem, conditional on the first-stage path. Prior(Chose Suboptimally) is constructed from reported beliefs using equation (3). Both prior and posterior beliefs are expressed in log-odds. Log-Likelihood(Path Move) is the empirical log-likelihood of a given path move in a particular sequential choice problem. Off Path is an indicator variable equal to one if the binary choice problem follows an off-path move and zero otherwise. Diff Max Payoff denotes the difference in maximum continuation payoffs across the on- and off-path binary choice problems. Payoffs are measured in pounds. Choice data are from part 2 of the choice-elicitation experiment and belief data from part 2 of the belief-elicitation experiment. Standard errors clustered at the participant level are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.

with selection neglect: participants recognise that off-path moves are informative, but underweight the extent to which they select less accurate decision-makers.

One reason for this underinference is that the relevant inference is cognitively demanding. Participants must condition on a low-probability path, infer what that path reveals about the decision-maker, and then apply that inference to a subsequent choice in the selected continuation problem. This resembles selection neglect and inference from selected samples (Jehiel, 2018; Esponda and Vespa, 2018; Barron et al., 2024), where, in our setting, the selected sample is the set of participants who have just made an off-path move. More broadly, it relates to failures of contingent reasoning, where participants underuse information that becomes relevant only after conditioning on particular states or histories (Martinez-Marquina et al., 2019; Niederle and Vespa, 2023; Esponda and Vespa, 2024; Calford and Cason, 2024). It is also consistent with evidence that complex information structures generate muted belief updating (Gonçalves et al., 2025; Ba et al., 2025).

Column (3) shows that participants rely more on path informativeness when the first-stage deviation involves a larger difference in foregone maximum payoffs. This is consistent with the evidence in Table 1: participants understand that more severe deviations are more informative. However, the level of updating remains far from the empirical benchmark. Regarding priors, the coefficient on prior log-odds is consistently close to, but significantly below, one. Participants therefore exhibit mild base-rate neglect in addition to substantial underinference from the path signal.

In short, the data suggest that participants understand that off-path moves are informative, but substantially underuse the information contained in those moves. Off-path histories generate strong selection in the choice data, and beliefs exhibit selection neglect.

## 7. Perfection Believers

We now study participants who assign zero, or nearly zero, probability to off-path moves. These participants are of particular interest because they allow us to study beliefs following events that are subjectively impossible or extremely unlikely. This is precisely the type of situation that motivates refinements of Nash equilibrium in dynamic games: off-path histories are not expected to occur, but sequential reasoning nevertheless requires beliefs and optimal continuation behaviour after such histories (Selten, 1975; Kreps and Wilson, 1982; Reny, 1992, 1995).<sup>11</sup>

In our data, 13% of participants in the belief experiment assign probability zero to off-path moves in every first-stage belief report in part 2. These participants report beliefs indicating that others will *never* move off path. We refer to participants who always assign probability zero to off-path moves as *perfection believers*. This contrasts with the median participant, whose average reported belief about off-path moves is 17%.

We split participants in the belief experiment into four groups. The first group consists of perfection believers, who always assign probability zero to off-path moves. The remaining groups consist of participants who, on average, assign probability between zero and 5%, between 5% and 10%, and above 10% to off-path moves. We then ask whether beliefs about suboptimal continuation choices differ across these groups when participants condition on on- versus off-path moves.

Figure 9 shows the corresponding path effects on beliefs. The figure reports the change in beliefs about suboptimal choices following on- and off-path moves, relative to one-shot settings, separately for the four groups. These effects correspond to the coefficients from the specification in column (4) of Table 1, estimated separately by group.

The data show two main patterns. First, perfection believers, who expect there to be no off-path moves, also do *not* expect more mistakes in subsequent choices relative to one-shot problems, even after an off-path move. Their conditional beliefs are close to what one would expect from a literal application of subgame perfection or sequential rationality in our setting: even after conditioning on an off-path move, they continue to expect nearly optimal continuation behaviour. In this sense, perfection believers do not revise their model of choice behaviour after being asked to condition on an event that their own first-stage beliefs treated as impossible.

This result is naturally related to belief revision after unforeseen events. Becker et al. (2026) study how people revise beliefs after events outside their ex ante model. They find that, after an unforeseen

<sup>11</sup>We do note that, in many games, there are also paths of play that cannot be rationalised not just by subgame perfect Nash equilibrium, but by *any* Nash equilibrium, posing conceptual problems already for players intent on adhering to a model of their opponents' behaviour consistent with observed choices.






Figure 9: Perfection Believers and Beliefs On- and Off-Path

*Notes:* This figure shows the change in beliefs about suboptimal choices in binary choice problems, expressed in percentage points, following on- and off-path moves relative to one-shot settings. Participants are grouped by their average reported probability of off-path moves: 0% (brown circles), between 0% and 5% (red diamonds), between 5% and 10% (orange triangles), and more than 10% (yellow crosses). The plotted values are coefficients from regressions corresponding to column (4) of [Table 1](#), estimated separately for each group. One-shot beliefs are from part 1 of the belief-elicitation experiment. Beliefs following on- and off-path moves are from the second stage of part 2 of the same experiment. Whiskers correspond to 95% confidence intervals using standard errors clustered at the participant level.

event occurs, belief revisions are broadly consistent with reverse Bayesianism: agents incorporate the newly recognised event whilst largely preserving the relative likelihoods of previously considered states ([Karni and Vierø, 2013](#)). Our setting differs because off-path moves are objectively part of the model, but perfection believers subjectively rule them out. Our data therefore speak to a related question: how do people reason after events to which they had assigned probability zero? The evidence suggests that perfection believers do not substantially revise their model of others' behaviour after conditioning on the subjectively impossible event.

Second, participants who assign positive probability to off-path moves are more willing to condition on the path. Those who assign probabilities close to zero behave similarly to perfection believers. By contrast, participants who assign probabilities between 5% and 10% to off-path moves report substantially higher beliefs about suboptimal choices off path. They therefore respond more strongly to the selection effect documented in [Section 6](#). The appendix provides a finer partition of first-stage beliefs. [Table 8](#) in [Appendix A.4](#) shows that participants assigning between 20% and 30% probability to off-path moves condition most strongly on off-path histories, with weaker conditioning among participants who assign both lower and higher probabilities.

These results connect to experimental and theoretical work on off-path beliefs and higher-order reasoning in dynamic games. In centipede and related games, players' behaviour after unexpected histories depends on what they infer from deviations about the opponent's type or preferences ([Wang, 2023](#); [Breitmoser et al., 2014](#); [Ho and Su, 2013](#); [Healy, 2024](#)). Our design removes strategic-



preference confounds and shows a related phenomenon in a simpler environment. Some participants treat off-path moves as informative evidence about future mistakes. Perfection believers, by contrast, appear to maintain a near-perfect model of behaviour even when asked to reason after a subjectively zero-probability event.

## 8. Response Times

The previous sections suggest that mistakes, even in our simple choice environment, may be related to cognitive frictions. In this section, we use response times as auxiliary evidence on those frictions. The analysis is exploratory, but it provides a useful behavioural marker: choices are slower when the relevant payoff comparisons are closer.

A large literature relates response times to preference intensity. In binary choice problems, models of sequential sampling and noisy evidence accumulation often predict that decisions are slower when the value difference between alternatives is smaller. The same environments also tend to generate noisier choices. This negative relationship between strength of preference and response time has been documented at least since [Mosteller and Nogee \(1951\)](#)<sup>12</sup> and is central in sequential-sampling models of economic choice (e.g., [Fehr and Rangel, 2011](#); [Fudenberg et al., 2018](#); [Alós-Ferrer et al., 2021](#)).

[Table 3a](#) reports regressions of log response time on payoff features. Column (1) focuses on one-shot binary choice problems. Columns (2) to (5) focus on sequential choice problems. In [Appendix A.5](#), we report the same regressions using response time in levels instead of its natural logarithm; the results are qualitatively consistent throughout.

In the simple one-shot problems, choices are fast and response times vary little with payoff features. This is consistent with the evidence in [Figure 5](#): one-shot choices are nearly always optimal, and payoff variation has only limited scope to affect behaviour.

The response-time patterns are stronger in the sequential problems. As we have seen in [Figure 6](#), first-stage choices are more sensitive to payoff features in the sequential environment. Response times move in the corresponding direction. Choices are slower when the difference in maximum or expected continuation payoffs is smaller. Thus, the same payoff comparison that predicts more off-path moves also predicts longer decision times. Response times are also slower when the difference in minimum continuation payoffs is smaller, consistent with the earlier finding that minimum-payoff comparisons affect first-stage choices.

These patterns indicate that response times track preference intensity. They do not, however, simply proxy for participant-level mistake propensity. Despite substantial between-participant heterogeneity in response times, participants who are slower on average are not systematically more likely to make one-shot mistakes or off-path moves. This is shown in [Table 10](#) in [Appendix A.5](#). Whilst

---

<sup>12</sup>Recent references include [Krajbich et al. \(2012\)](#), [Spiliopoulos and Ortmann \(2018\)](#), [Clithero \(2018\)](#), and [Konovalov and Krajbich \(2019\)](#).

it may seem at odds with the posited heterogeneity in cognitive frictions, the absence of a monotone relationship between response time and individual mistake propensity is not inconsistent with cognitive models of choice. In sequential-sampling models, response time reflects both preference intensity and the decision-maker's ability, and it will generically be non-monotone in ability—cf. [Gonçalves \(2026\)](#).

We also examine response times in the belief-elicitation experiment. [Table 3b](#) reports regressions analogous to those in [Table 3a](#), but using the time taken to report beliefs rather than the time taken to make choices. Belief reports are faster when payoff differences are larger. This suggests that similar cognitive forces affect the assessment of others' choices: belief participants are faster when the payoff comparison makes the predicted choice more straightforward.

The participant-level relationship differs from the choice experiment. In the belief data, longer response times are associated with higher reported beliefs about others' mistake propensities, as shown in [Table 10](#) in [Appendix A.5](#). One interpretation is that believing others choose optimally is a relatively instinctive response, whereas forming beliefs about when and why others make mistakes requires more deliberation. This interpretation is consistent with evidence that response times distinguish more instinctive from more contemplative responses in economic decisions ([Rubinstein, 2016](#); [Spiliopoulos and Ortmann, 2018](#)).

|                 | Log(Time)           |                      |                      |                      |                      |
|-----------------|---------------------|----------------------|----------------------|----------------------|----------------------|
|                 | One-Shot            | Sequential           |                      |                      |                      |
|                 | (1)                 | (2)                  | (3)                  | (4)                  | (5)                  |
| Payoff Spread   | -0.001<br>(0.020)   | –                    | –                    | –                    | –                    |
| Payoff Level    | 0.041*<br>(0.024)   | –                    | –                    | –                    | –                    |
| Diff Max Payoff | –                   | -0.101***<br>(0.009) | -0.077***<br>(0.010) | –                    | –                    |
| Diff Exp Payoff | –                   | –                    | –                    | -0.102***<br>(0.009) | -0.078***<br>(0.010) |
| Diff Min Payoff | –                   | –                    | -0.047***<br>(0.011) | –                    | -0.046***<br>(0.011) |
| Intercept       | 1.097***<br>(0.036) | 1.739***<br>(0.020)  | 1.733***<br>(0.020)  | 1.739***<br>(0.020)  | 1.733***<br>(0.020)  |
| R <sup>2</sup>  | 0.00                | 0.01                 | 0.01                 | 0.01                 | 0.01                 |
| N               | 2025                | 5912                 | 5912                 | 5912                 | 5912                 |

(a) Choices

|                 | Log(Time)            |                      |                      |                      |                      |
|-----------------|----------------------|----------------------|----------------------|----------------------|----------------------|
|                 | One-Shot             | Sequential           |                      |                      |                      |
|                 | (1)                  | (2)                  | (3)                  | (4)                  | (5)                  |
| Payoff Spread   | -0.099***<br>(0.027) | –                    | –                    | –                    | –                    |
| Payoff Level    | 0.078**<br>(0.035)   | –                    | –                    | –                    | –                    |
| Diff Max Payoff | –                    | -0.108***<br>(0.014) | -0.085***<br>(0.017) | –                    | –                    |
| Diff Exp Payoff | –                    | –                    | –                    | -0.116***<br>(0.015) | -0.094***<br>(0.018) |
| Diff Min Payoff | –                    | –                    | -0.044***<br>(0.017) | –                    | -0.037**<br>(0.017)  |
| Intercept       | 1.923***<br>(0.051)  | 3.009***<br>(0.033)  | 3.003***<br>(0.033)  | 3.010***<br>(0.033)  | 3.004***<br>(0.033)  |
| R <sup>2</sup>  | 0.00                 | 0.01                 | 0.01                 | 0.01                 | 0.01                 |
| N               | 2025                 | 5905                 | 5905                 | 5905                 | 5905                 |

(b) Beliefs

Table 3: Response Times and Payoffs

*Notes:* This table shows how response times in choices (subtable (a)) and belief reports (subtable (b)) vary with payoff features in one-shot binary choice problems, column (1), and in sequential choice problems, columns (2) to (5). Payoff Spread is the difference between the two payoffs in a binary choice problem. Payoff Level is the minimum payoff in the binary choice problem. Diff Max/Exp/Min Payoff denotes the difference in maximum/expected/minimum continuation payoffs between the on- and off-path binary choice problems in sequential choice problems. Payoffs are measured in pounds. Time is response time in seconds and Log(Time) is its natural logarithm. The data comprise one-shot binary choice problems in part 1 and sequential choice problems in part 2 of the choice- and belief-elicitation experiments. Standard errors clustered at the participant level are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.

## 9. Conclusion

Off-path behaviour is usually treated as a theoretical object: something that must be specified for equilibrium reasoning, but that is not expected to occur. This paper studies off-path behaviour and beliefs empirically. We use a simple experimental design in which off-path moves are directly identified, mistakes are well defined, and beliefs about choices following on- and off-path moves can be elicited separately.

The data show that off-path moves are rare, but not uninformative. Participants choose suboptimally in 2.4% of one-shot problems and move off path in 5.7% of first-stage choices. Belief participants also recognise that suboptimal choices and off-path events are unlikely, although they overestimate their frequency.

Yet, when an off-path move occurs, it is highly diagnostic of subsequent behaviour. Participants who move off path are substantially more likely to make later suboptimal choices, especially when the first-stage deviation is more costly. In this sense, an off-path move is not merely an isolated tremble.

The belief data reveal a more subtle pattern. Participants understand that off-path moves contain information. They assign higher probabilities to subsequent mistakes after off-path moves and lower probabilities after on-path moves. However, they underuse the information contained in the path. The choice data exhibit strong selection into off-path histories, whereas beliefs only partially incorporate that selection.

This gap between off-path behaviour and beliefs is largest among participants who most strongly believe in perfection. Perfection believers assign zero probability to off-path moves and, even when asked to condition on such moves, continue to report beliefs close to a perfect-behaviour model.

These findings point to a distinction between normative and positive uses of off-path reasoning. Sequential rationality and perfection provide disciplined ways to restrict behaviour after histories that should not occur. Our results suggest that people do not ignore such histories when they occur, but neither do they treat them as fully informative signals. They recognise payoff-dependent mistakes and they recognise selection, but they do not combine these pieces of information as strongly as the data warrant. The empirical challenge is therefore not simply that people fail to reason off path. It is that off-path reasoning is selective: people condition on the unexpected, but too weakly, and least of all when they were most committed to the event not occurring.

In our design, we can isolate off-path moves that are unambiguous mistakes and that observers understand as mistakes. An off-path move is therefore informative about a player's propensity to make mistakes, and participants' beliefs respond to this, though less strongly than what the choice data implies. In richer settings, other factors contribute to differences between an analyst and players as to whether a move is on or off path. In centipede games, for instance, first-stage decisions may or may not be interpreted by players as revealing the first-mover's social preferences. An analyst then has to disentangle whether a move was considered to be an off-path mistake as in our setting,

or is instead interpreted by other players as simply reflecting (social) preferences over outcomes. In contrast, moves that may not have been considered off-path by the analyst, may be deemed to be so by players since they expected others to play according to a different equilibrium. Future work can reintroduce these motives and ask how observers separate them from the inference we isolate here.

## References

- Agranov, M., and P. Detkova.** 2025. “How People Think Others Update Beliefs: An Experiment.” *Working Paper*.
- Agranov, M., and P. Reshidi.** 2026. “Disentangling Suboptimal Updating: Task Difficulty, Structure, and Sequencing.” *Working Paper*.
- Almond, D., and D. Martin.** 2024. “Rational Inattention in Games: Experimental Evidence.” *Experimental Economics* 27 715–742. 10.1007/s10683-024-09843-z.
- Alós-Ferrer, C., E. Fehr, and N. Netzer.** 2021. “Time Will Tell: Recovering Preferences When Choices Are Noisy.” *Journal of Political Economy* 129 (6): 1828–1877. 10.1086/713732.
- Andreoni, J., and E. Blanchard.** 2006. “Testing Subgame Perfection Apart from Fairness in Ultimatum Games.” *Experimental Economics* 9 (4): 307–321. 10.1007/s10683-006-0064-7.
- Aoyagi, M., G. R. Fréchette, and S. Yuksel.** 2024. “Beliefs in Repeated Games: An Experiment.” *American Economic Review* 114 (12): 3944–3975. 10.1257/aer.20220639.
- Azrieli, Y., C. P. Chambers, and P. J. Healy.** 2018. “Incentives in Experiments: A Theoretical Analysis.” *Journal of Political Economy* 126 (4): 1472–1503. 10.1086/698136.
- Ba, C., J. A. Bohren, and A. Imas.** 2025. “Over- and Underreaction to Information: Belief Updating with Cognitive Constraints.” *Working Paper*.
- Baliga, S., and T. Sjöström.** 2004. “Arms Races and Negotiations.” *Review of Economic Studies* 71 (2): 351–369. 10.1111/0034-6527.00287.
- Barron, K., S. Huck, and P. Jehiel.** 2024. “Everyday Econometricians: Selection Neglect and Overoptimism When Learning from Others.” *American Economic Journal: Microeconomics* 16 (3): 162–198. 10.1257/mic.20200030.
- Basu, K.** 1990. “On the Non-Existence of a Rationality Definition for Extensive Games.” *International Journal of Game Theory* 19 (1): 33–44. 10.1007/BF01753706.
- Battigalli, P., and M. Siniscalchi.** 2002. “Strong Belief and Forward Induction Reasoning.” *Journal of Economic Theory* 106 (2): 356–391. 10.1006/jeth.2001.2942.
- Beard, T. R., and R. O. Beil.** 1994. “Do People Rely on the Self-Interested Maximization of Others? An Experimental Test.” *Management Science* 40 (2): 252–262. 10.1287/mnsc.40.2.252.
- Becker, C. K., T. Melkonyan, E. Proto, A. Sofianos, and S. T. Trautmann.** 2026. “Revising Beliefs in Light of Unforeseen Events.” *Journal of the European Economic Association* 24 (2): 661–700. 10.1093/jeea/jvaf038.
- Benjamin, D. J.** 2019. “Errors in Probabilistic Reasoning and Judgment Biases.” In *Handbook of Behavioral Economics: Applications and Foundations 1*, edited by Bernheim, B. D., S. DellaVigna, and D. Laibson Volume 2. 69–186, North-Holland, . 10.1016/bs.hesbe.2018.11.002.
- Binmore, K.** 1987. “Modeling Rational Players: Part I.” *Economics and Philosophy* 3 (2): 179–214. 10.1017/S0266267100002893.

- Block, H. D., and J. Marschak.** 1960. “Random Orderings And Stochastic Theories of Responses.” In *Contributions to Probability and Statistics*, edited by Olkin, I., S. Ghurye, W. Hoeffding, W. Madow, and H. Mann 97–132, Stanford University Press.
- Breitmoser, Y., J. H. W. Tan, and D. J. Zizzo.** 2014. “On the Beliefs Off the Path: Equilibrium Refinement Due to Quantal Response and Level- $k$ .” *Games and Economic Behavior* 86 102–125. 10.1016/j.geb.2014.03.002.
- Calford, E. M., and T. N. Cason.** 2024. “Contingent Reasoning and Dynamic Public Goods Provision.” *American Economic Journal: Microeconomics* 16 (2): 236–266. 10.1257/mic.20220111.
- Camerer, C. F., T.-H. Ho, and J.-K. Chong.** 2004. “A Cognitive Hierarchy Model of Games.” *The Quarterly Journal of Economics* 119 (3): 861–898. 10.1162/0033553041502225.
- Caplin, A.** 2016. “Measuring and Modeling Attention.” *Annual Review of Economics* 8 (1): 379–403. 10.1146/annurev-economics-080315-015417.
- Caplin, A., and M. Dean.** 2015. “Revealed Preference, Rational Inattention, and Costly Information Acquisition.” *American Economic Review* 105 (7): 2183–2203. 10.1257/aer.20140117.
- Caplin, A., M. Dean, and J. Leahy.** 2022. “Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy.” *Journal of Political Economy* 130 (6): 1–40. 10.1086/719276.
- Chakraborty, A., and C. Kendall.** 2026. “Noisy Foresight.” *Journal of the European Economic Association* jvag013 1–38. 10.1093/jeea/jvag013.
- Chan, K.** 2026. “An Axiomatic Model and Test of Belief Updating: Bayes, Grether (1980), and Preservation of MLRP.” *Working Paper*.
- Charness, G., U. Gneezy, and B. Halladay.** 2021a. “Experimental methods: Pay one or pay all.” *Journal of Economic Behavior & Organization* 131 (A): 141–150. 10.1016/j.jebo.2016.08.010.
- Charness, G., U. Gneezy, and V. Rasocha.** 2021b. “Experimental methods: Eliciting beliefs.” *Journal of Economic Behavior & Organization* 189 (2021): 234–256. 10.1016/j.jebo.2021.06.032.
- Chassang, S., and G. Padró-i-Miquel.** 2010. “Conflict and Deterrence Under Strategic Risk.” *Quarterly Journal of Economics* 125 (4): 1821–1858. 10.1162/qjec.2010.125.4.1821.
- Cho, I.-K., and D. M. Kreps.** 1987. “Signaling Games and Stable Equilibria.” *The Quarterly Journal of Economics* 102 (2): 179–221. 10.2307/1885060.
- Clithero, J.** 2018. “Improving out-of-sample predictions using response times and a model of the decision process.” *Journal of Economic Behavior and Organization* 148 344–375. 10.1016/j.jebo.2018.02.007.
- Costa-Gomes, M. A., and G. Weizsäcker.** 2008. “Stated Beliefs and Play in Normal-Form Games.” *The Review of Economic Studies* 75 (3): 729–762. 10.1111/j.1467-937X.2008.00498.x.
- Cripps, M. W.** 2025. “Divisible Updating.” *Working Paper*.
- Danz, D., D. Fehr, and D. Kübler.** 2012. “Information and Beliefs in a Repeated Normal-Form Game.” *Experimental Economics* 15 (4): 622–640. 10.1007/s10683-012-9317-9.
- Danz, D., L. Vesterlund, and A. J. Wilson.** 2022. “Belief Elicitation and Behavioral Incentive Compatibility.” *American Economic Review* 112 (9): 2851–83. 10.1257/aer.20201248.
- Enke, B., and T. Graeber.** 2023. “Cognitive Uncertainty.” *Quarterly Journal of Economics* 138 (4): 2021–2067. 1093/qje/qjad025.
- Enke, B., T. Graeber, R. Oprea, and J. Yang.** 2026. “Behavioral Attenuation.” *Working Paper*.
- Esponda, I., and E. Vespa.** 2018. “Endogenous sample selection: A laboratory study.” *Quantitative*

*Economics* 9 183–216. 10.3982/QE650.

**Esponda, I., and E. Vespa.** 2024. “Contingent Thinking and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory.” *Review of Economic Studies* 91 (5): 2806–2831. 10.1093/restud/rdad102.

**Esteban-Casanelles, T., and D. Gonçalves.** 2026. “Incentives and Strategic Behaviour.” *Experimental Economics* Forthcoming.

**Fehr, E., and A. Rangel.** 2011. “Neuroeconomic Foundations of Economic Choice—Recent Advances.” *Journal of Economic Perspectives* 25 (4): 3–30. 10.1257/jep.25.4.3.

**Fehr, E., and K. M. Schmidt.** 1999. “A Theory of Fairness, Competition, and Cooperation.” *Quarterly Journal of Economics* 114 (3): 817–868. 10.1162/003355399556151.

**Fershtman, C., and A. Pakes.** 2012. “Dynamic Games with Asymmetric Information: A Framework for Empirical Work.” *Quarterly Journal of Economics* 127 (4): 1611–1661. 10.1093/qje/qjs025.

**Fong, M.-J., P.-H. Lin, and T. R. Palfrey.** 2025. “Cursed Sequential Equilibrium.” *American Economic Review* 115 (8): 2616–2658. 10.1257/aer.20231130.

**Friedman, E., and J. Ward.** 2022. “Stochastic Choice and Noisy Beliefs in Games: an Experiment.” *Working Paper* 1–77. 10.2139/ssrn.4190338.

**Fudenberg, D., P. Strack, and T. Strzalecki.** 2018. “Speed, Accuracy, and the Optimal Timing of Choices.” *American Economic Review* 108 (12): 3651–3684. 10.1257/aer.20150742.

**Fudenberg, D., R. Iijima, and T. Strzalecki.** 2015. “Stochastic Choice and Revealed Perturbed Utility.” *Econometrica* 83 (6): 2371–2409. 10.3982/ECTA12660.

**Gilboa, I., and D. Schmeidler.** 1989. “Maxmin expected utility with non-unique prior.” *Journal of Mathematical Economics* 18 (2): 141–153. 10.1016/0304-4068(89)90018-9.

**Goeree, J. K., and C. A. Holt.** 2001. “Ten Little Treasures of Game Theory and Ten Intuitive Contradictions.” *American Economic Review* 91 (5): 1402–1422. 10.1257/aer.91.5.1402.

**Gonçalves, D.** 2026. “Speed, Accuracy, and Complexity.” *Working Paper*.

**Gonçalves, D., J. Libgober, and J. Willis.** 2025. “Retractions: Updating from Complex Information.” *The Review of Economic Studies* 93 (1): 476–516. 10.1093/restud/rdaf032.

**Grether, D. M.** 1980. “Bayes Rule as a Descriptive Model: The Representativeness Heuristic.” *The Quarterly Journal of Economics* 95 (3): 537–557. 10.2307/1885092.

**Grether, D. M.** 1992. “Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence.” *Journal of Economic Behavior & Organization* 17 (1): 31–57. 10.1016/0167-2681(92)90078-P.

**Gul, F., and W. Pesendorfer.** 2001. “Temptation and Self-Control.” *Econometrica* 69 (6): 1403–1435. 10.1111/1468-0262.00252.

**Healy, P. J.** 2024. “Epistemic Experiments: Utilities, Beliefs, and Irrational Play.” Working paper, The Ohio State University.

**Healy, P. J., and J. Kagel.** 2025. “Testing Elicitation Mechanisms via Team Chat.” *Working Paper*.

**Healy, P. J., and G. Leo.** 2025. “Chapter 3 - Belief elicitation: a user’s guide.” In *Handbook of Experimental Methodology*, edited by Snowberg, E., and L. Yariv Volume 1. 81–162, North-Holland, . 10.1016/bs.hbem.2025.09.002.

**Ho, T.-H., and X. Su.** 2013. “A Dynamic Level- $k$  Model in Sequential Games.” *Management Science* 59 (1): 33–51. 10.1287/mnsc.1120.1645.



- Hossain, T., and R. Okui.** 2013. "The Binarized Scoring Rule." *Review of Economic Studies* 80 (3): 984–1001. 10.1093/restud/rdt006.
- Huck, S., and G. Weizsäcker.** 2002. "Do Players Correctly Estimate What Others Do? Evidence of Conservatism in Beliefs." *Journal of Economic Behavior & Organization* 47 (1): 71–85. 10.1016/S0167-4870(99)00031-8.
- Hyndman, K., E. Y. Özbay, A. Schotter, and W. Ehrblatt.** 2012. "Belief Formation: An Experiment with Outside Observers." *Experimental Economics* 15 (1): 176–203. 10.1007/s10683-011-9296-2.
- Jehiel, P.** 2018. "Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism." *American Economic Review* 108 (6): 1582–1597. 10.1257/aer.20161696.
- Karni, E., and M.-L. Vierø.** 2013. "Reverse Bayesianism": A Choice-Based Theory of Growing Awareness." *American Economic Review* 103 (7): 2790–2810. 10.1257/aer.103.7.2790.
- Ke, S.** 2019. "Boundedly rational backward induction." *Theoretical Economics* 14 103–134. 10.3982/TE2402.
- Konovalov, A., and I. Krajbich.** 2019. "Revealed strength of preference: Inference from response times." *Judgment and Decision Making* 14 (4): 381–394. 10.1017/S1930297500006082.
- Krajbich, I., D. Lu, C. Camerer, and A. Rangel.** 2012. "The attentional drift-diffusion model extends to simple purchasing decisions." *Frontiers in Psychology* 3 1–18. 10.3389/fpsyg.2012.00193.
- Kreps, D. M.** 2023. *Microeconomic Foundations II: Imperfect Competition, Information, and Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Kreps, D. M., and R. Wilson.** 1982. "Sequential Equilibria." *Econometrica* 50 (4): 863–894. 10.2307/1912767.
- Levitt, S. D., J. A. List, and S. E. Sadoff.** 2011. "Checkmate: Exploring Backward Induction among Chess Players." *American Economic Review* 101 (2): 975–990. 10.1257/aer.101.2.975.
- Lin, P.-H., and T. R. Palfrey.** 2024. "Cognitive Hierarchies for Games in Extensive Form." *Journal of Economic Theory* 220 105871. 10.1016/j.jet.2024.105871.
- Luce, R. D.** 1959. *Individual Choice Behavior: A Theoretical Analysis*. New York, NY: John Wiley & Sons.
- Martinez-Marquina, A., M. Niederle, and E. Vespa.** 2019. "Failures in Contingent Reasoning? The Role of Uncertainty." *American Economic Review* 109 (10): 3437–3474. 10.1257/aer.20171764.
- Mas-Colell, A., M. D. Whinston, and J. R. Green.** 1995. *Microeconomic Theory*. Oxford University Press.
- Mattsson, L.-G., and J. W. Weibull.** 2002. "Probabilistic Choice and Procedural Rationality." *Games and Economic Behavior* 41 (1): 61–78. 10.1016/S0899-8256(02)00014-3.
- Matějka, F., and A. McKay.** 2015. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." *American Economic Review* 105 (1): 272–298. 10.1257/aer.20130047.
- McFadden, D.** 1974. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics*, edited by Zarembka, P. 105–142, New York, NY: Academic Press.
- McKelvey, R. D., and T. R. Palfrey.** 1998. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics* 1 9–41. 10.1023/A:1009905800005.
- McKelvey, R. D., and T. R. Palfrey.** 1992. "An Experimental Study of the Centipede Game." *Econometrica* 60 (4): 803–836. 10.2307/2951567.
- McKelvey, R. D., and T. R. Palfrey.** 1995. "Quantal Response Equilibria for Normal Form Games."

- Games and Economic Behavior* 10 (1): 6–38. 10.1006/game.1995.1023.
- Mosteller, F., and P. Nogee.** 1951. “An Experimental Measurement of Utility.” *Journal of Political Economy* 59 (5): 371–404. 10.1086/257106.
- Nagel, R.** 1995. “Unraveling in Guessing Games: An Experimental Study.” *American Economic Review* 85 (5): 1313–1326.
- Niederle, M., and E. Vespa.** 2023. “Cognitive Limitations: Failures of Contingent Thinking.” *Annual Review of Economics* 15 (2023): 307–328. 10.1146/annurev-economics-091622-124733.
- Nocke, V., and M. D. Whinston.** 2010. “Dynamic Merger Review.” *Journal of Political Economy* 118 (6): 1201–1251. 10.1086/658161.
- Nyarko, Y., and A. Schotter.** 2002. “An Experimental Study of Belief Learning Using Elicited Beliefs.” *Econometrica* 70 (3): 971–1005. 10.1111/1468-0262.00316.
- Palacios-Huerta, I., and O. Volij.** 2009. “Field Centipedes.” *American Economic Review* 99 (4): 1619–35. 10.1257/aer.99.4.1619.
- Palfrey, T. R., and S. W. Wang.** 2009. “On eliciting beliefs in strategic games.” *Journal of Economic Behavior and Organization* 71 (2): 98–109. 10.1016/j.jebo.2009.02.009.
- Reny, P. J.** 1992. “Backward Induction, Normal Form Perfection and Explicable Equilibria.” *Econometrica* 60 (3): 627–649. 10.2307/2951586.
- Reny, P. J.** 1995. “Rational Behaviour in Extensive-Form Games.” *The Canadian Journal of Economics / Revue canadienne d’Economie* 28 (1): 1–16. 10.2307/136020.
- Rosenthal, R.** 1981. “Games of Perfect Information, Predatory Pricing, and the Chain Store Paradox.” *Journal of Economic Theory* 24 92–100. 10.1016/0022-0531(81)90018-1.
- Rubinstein, A.** 2016. “A Typology of Players: Between Instinctive and Contemplative.” *Quarterly Journal of Economics* 131 (2): 859–890. 10.1093/qje/qjw008.
- Ryan, S. P.** 2012. “The Costs of Environmental Regulation in a Concentrated Industry.” *Econometrica* 80 (3): 1019–1061. 10.3982/ECTA6750.
- Schlag, K., J. Tremewan, and J. J. van der Weele.** 2015. “A penny for your thoughts: A survey of methods for eliciting beliefs.” *Experimental Economics* 18 (3): 457–490. 10.1007/s10683-014-9416-x.
- Schotter, A., and I. Trevino.** 2014. “Belief Elicitation in the Laboratory.” *Annual Review of Economics* 6 103–128. 10.1146/annurev-economics-080213-040927.
- Selten, R.** 1975. “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games.” *International Journal of Game Theory* 4 (1): 25–55. 10.1007/BF01766400.
- Shubatt, C., and J. Yang.** 2026. “Tradeoffs and Comparison Complexity.” *Working Paper*.
- Spiliopoulos, L., and A. Ortmann.** 2018. “The BCD of response time analysis in experimental economics.” *Experimental Economics* 21 383–433. 10.1007/s10683-017-9528-1.
- Stahl, D. O., and P. W. Wilson.** 1995. “On Players’ Models of Other Players: Theory and Experimental Evidence.” *Games and Economic Behavior* 10 (1): 218–254. 10.1006/game.1995.1031.
- Strotz, R. H.** 1955. “Myopia and Inconsistency in Dynamic Utility Maximization.” *Review of Economic Studies* 23 (3): 165–180. 10.2307/2295722.
- Strzalecki, T.** 2023. *Stochastic Choice Theory*. Princeton Lectures in Economics, Princeton, NJ: Princeton University Press.
- Toussaert, S.** 2018. “Eliciting Temptation and Self-Control Through Menu Choices: A Lab Experi-

ment.” *Econometrica* 86 (3): 859–889. 10.3982/ECTA14172.

**Wang, Y.** 2023. “Belief and Higher-Order Belief in the Centipede Games: An Experimental Investigation.” *Pacific Economic Review* 28 (1): 27–73. 10.1111/1468-0106.12404.

**Wilson, A. J., and E. Vespa.** 2016. “Paired-Uniform Scoring: Implementing a Binarized Scoring Rule with Non-Mathematical Language.” *Working Paper*.

## A. Supporting Tables and Figures

### A.1. Choices and Beliefs in One-Shot Settings

In [Section 3](#), we show that payoff spreads affect choices and beliefs in one-shot settings, whereas payoff levels do not. [Table 4](#) reports the corresponding regressions. Columns (1) and (2) show the effects of payoff spreads and payoff levels on choices and beliefs, respectively. Column (3) reports the effect on belief bias, defined as the absolute difference between the reported belief about suboptimal choice and the empirical frequency of suboptimal choice for that decision problem. The intercept shows substantial overestimation of the baseline mistake propensity. In addition, although payoff spreads are informative about choices, participants overinfer from them: they expect a stronger decrease in mistakes as the payoff spread increases than is observed in the choice data.

|                | Chose Suboptimally (One Shot) |                       |                      |
|----------------|-------------------------------|-----------------------|----------------------|
|                | Choice                        | Belief                | Belief Bias          |
|                | (1)                           | (2)                   | (3)                  |
| Payoff Spread  | -2.125**<br>(1.018)           | -10.058***<br>(0.801) | -8.984***<br>(0.719) |
| Payoff Level   | 0.036<br>(0.567)              | 0.385<br>(0.572)      | 0.302<br>(0.553)     |
| Intercept      | 3.683***<br>(1.218)           | 20.186***<br>(1.303)  | 19.057***<br>(1.153) |
| R <sup>2</sup> | 0.00                          | 0.06                  | 0.06                 |
| N              | 2025                          | 2025                  | 2025                 |

Table 4: Choices and Beliefs in One-Shot Binary Choice Problems

*Notes:* This table shows the effect of the payoff spread and the payoff level of one-shot binary choice problems on the frequency of suboptimal choices, column (1), reported beliefs about suboptimal choices, column (2), and belief bias, column (3). Belief bias is the absolute difference between reported beliefs and empirical choice frequencies. Outcomes are expressed in percentage points. Payoffs are measured in pounds. Standard errors clustered at the participant level are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.

### A.2. The Off-Path

[Table 5](#) and [Table 6](#) report regressions of off-path moves and beliefs about off-path moves on payoff features of the sequential choice problems. [Table 5](#) focuses on the difference in maximum and expected continuation payoffs between the on-path and off-path continuation problems. [Table 6](#) adds the difference in minimum continuation payoffs. As in the one-shot problems, reported beliefs overestimate the baseline frequency of mistakes, as indicated by the intercepts. Beliefs are also more responsive to payoff variation than choices are. That is, participants believe payoff features of the sequential choice problem to be more diagnostic than they empirically are. We do not report regressions including both maximum and expected continuation payoffs because the two variables are strongly collinear.

|                 | Chose Off-Path       |                      |                       |                       |                       |                       |
|-----------------|----------------------|----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
|                 | Choice               |                      | Belief                |                       | Belief Bias           |                       |
|                 | (1)                  | (2)                  | (3)                   | (4)                   | (5)                   | (6)                   |
| Diff Max Payoff | -8.655***<br>(0.942) | –                    | -14.679***<br>(0.834) | –                     | -10.290***<br>(0.573) | –                     |
| Diff Exp Payoff | –                    | -8.775***<br>(0.954) | –                     | -16.169***<br>(0.912) | –                     | -11.331***<br>(0.623) |
| Intercept       | 12.299***<br>(1.063) | 12.344***<br>(1.068) | 28.116***<br>(1.191)  | 28.585***<br>(1.206)  | 22.853***<br>(0.694)  | 23.180***<br>(0.695)  |
| R <sup>2</sup>  | 0.03                 | 0.03                 | 0.09                  | 0.10                  | 0.07                  | 0.08                  |
| N               | 5912                 | 5912                 | 11810                 | 11810                 | 11810                 | 11810                 |

Table 5: Choices and Beliefs about Off-Path Moves

*Notes:* This table shows the effect of the difference in maximum and expected continuation payoffs between the on-path and off-path continuation binary choice problems on the frequency of off-path moves, columns (1)–(2), reported beliefs about off-path moves, columns (3)–(4), and belief bias, columns (5)–(6). Payoff differences are measured in pounds. Outcomes are measured in percentage points. Belief bias is the absolute difference between reported beliefs and empirical choice frequencies. Standard errors clustered at the participant level are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.

|                 | Chose Off-Path       |                      |                       |                       |                      |                      |
|-----------------|----------------------|----------------------|-----------------------|-----------------------|----------------------|----------------------|
|                 | Choice               |                      | Belief                |                       | Belief Bias          |                      |
|                 | (1)                  | (2)                  | (3)                   | (4)                   | (5)                  | (6)                  |
| Diff Min Payoff | -9.140***<br>(1.050) | -9.109***<br>(1.049) | -13.654***<br>(0.930) | -13.037***<br>(0.920) | -9.571***<br>(0.708) | -9.140***<br>(0.710) |
| Diff Max Payoff | -3.823***<br>(0.556) | –                    | -7.420***<br>(0.564)  | –                     | -5.202***<br>(0.445) | –                    |
| Diff Exp Payoff | –                    | -3.849***<br>(0.562) | –                     | -8.134***<br>(0.611)  | –                    | -5.699***<br>(0.481) |
| Intercept       | 11.125***<br>(0.954) | 11.117***<br>(0.953) | 26.297***<br>(1.126)  | 26.339***<br>(1.125)  | 21.578***<br>(0.682) | 21.605***<br>(0.681) |
| R <sup>2</sup>  | 0.06                 | 0.06                 | 0.16                  | 0.16                  | 0.13                 | 0.13                 |
| N               | 5912                 | 5912                 | 11810                 | 11810                 | 11810                | 11810                |

Table 6: Choices and Beliefs about Off-Path Moves; Minimum Payoffs

*Notes:* This table shows the effect of the difference in minimum continuation payoffs between the on-path and off-path continuation binary choice problems, controlling for the difference in maximum or expected continuation payoffs, on the frequency of off-path moves, columns (1)–(2), reported beliefs about off-path moves, columns (3)–(4), and belief bias, columns (5)–(6). Payoffs are measured in pounds. Outcomes are expressed in percentage points. Belief bias is the absolute difference between reported beliefs and empirical choice frequencies. Standard errors clustered at the participant level are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.

### A.3. Selection Effects and Selection Neglect

In [Section 6](#), we estimate a Grether-style updating regression in log-odds. Here, we report the corresponding specification in levels. The results are qualitatively similar.

|                                                          | Posterior(Chose Suboptimally) |                      |                      |
|----------------------------------------------------------|-------------------------------|----------------------|----------------------|
|                                                          | (1)                           | (2)                  | (3)                  |
| Prior(Chose Suboptimally)                                | 0.910***<br>(0.013)           | 0.913***<br>(0.013)  | 0.917***<br>(0.014)  |
| $\Delta$ Frequency(Path Move)                            | 0.122***<br>(0.009)           | 1.778***<br>(0.143)  | 1.775***<br>(0.144)  |
| $\Delta$ Frequency(Chose Suboptimally) x Off Path        | –                             | -1.673***<br>(0.139) | -1.801***<br>(0.147) |
| $\Delta$ Frequency(Chose Suboptimally) x Diff Max Payoff | –                             | –                    | 0.127***<br>(0.015)  |
| Intercept                                                | 1.523***<br>(0.223)           | 2.234***<br>(0.243)  | 2.289***<br>(0.242)  |
| R <sup>2</sup>                                           | 0.73                          | 0.73                 | 0.74                 |
| N                                                        | 11810                         | 11810                | 11810                |

Table 7: Belief Updating from Path Moves; Levels

*Notes:* This table reports estimates of the levels counterpart to the belief-updating regression in equation (2). Column (1) reports the baseline specification, and columns (2) and (3) add interaction terms. Posterior(Chose Suboptimally) is the reported belief about another participant’s second-stage choice in the sequential choice problem, conditional on the first-stage path. Prior(Chose Suboptimally) is constructed from reported beliefs using equation (3). Both prior and posterior beliefs are measured in levels.  $\Delta$  Frequency(Chose Suboptimally) is the empirical difference, for a given sequential choice problem, between the frequency of suboptimal choices conditional on the path move and the unconditional frequency of suboptimal choices. It is the levels counterpart of Log-Likelihood(Path Move) in Table 2: were these in log-odds instead of level frequencies, it would correspond to the log-likelihood of the path move as in Table 2. Off Path is an indicator variable equal to one if the binary choice problem follows an off-path move and zero otherwise. Diff Max Payoff denotes the difference in maximum continuation payoffs across the on- and off-path binary choice problems. Payoffs are measured in pounds. Choice data are from part 2 of the choice-elicitation experiment and belief data from part 2 of the belief-elicitation experiment. Standard errors clustered at the participant level are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.

#### A.4. Perfection Believers

| Avg Belief Off-Path: | Belief Chose Suboptimally |                     |                     |                      |                      |                      |                      |
|----------------------|---------------------------|---------------------|---------------------|----------------------|----------------------|----------------------|----------------------|
|                      | 0%                        | (0%, 5%]            | (5%, 10%]           | (10%, 20%]           | (20%, 30%]           | (30%, 40%]           | >40%                 |
|                      | (1)                       | (2)                 | (3)                 | (4)                  | (5)                  | (6)                  | (7)                  |
| Off Path             | -0.082<br>(0.589)         | 0.915<br>(1.188)    | 4.462***<br>(1.132) | 4.201***<br>(0.899)  | 7.781***<br>(0.815)  | 6.346***<br>(1.987)  | 3.209*<br>(1.764)    |
| On Path              | -0.490<br>(0.491)         | -1.102<br>(0.872)   | -1.125*<br>(0.653)  | -2.373***<br>(0.757) | -1.715*<br>(1.011)   | -0.274<br>(2.253)    | -0.307<br>(1.802)    |
| Intercept            | 0.490<br>(0.491)          | 2.232***<br>(0.848) | 4.060***<br>(0.967) | 9.883***<br>(1.144)  | 22.118***<br>(1.234) | 30.565***<br>(2.253) | 42.168***<br>(2.862) |
| N Participants       | 29                        | 23                  | 24                  | 60                   | 52                   | 24                   | 13                   |
| R <sup>2</sup>       | 0.00                      | 0.01                | 0.06                | 0.04                 | 0.08                 | 0.03                 | 0.01                 |
| N                    | 1783                      | 1419                | 1476                | 3682                 | 3198                 | 1474                 | 803                  |

(a)

| Avg Belief Off-Path:       | Belief Chose Suboptimally |                     |                      |                      |                      |                      |                      |
|----------------------------|---------------------------|---------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
|                            | 0%                        | (0%, 5%]            | (5%, 10%]            | (10%, 20%]           | (20%, 30%]           | (30%, 40%]           | >40%                 |
|                            | (1)                       | (2)                 | (3)                  | (4)                  | (5)                  | (6)                  | (7)                  |
| Off Path                   | -0.181<br>(0.557)         | -0.727<br>(1.392)   | 1.868**<br>(0.792)   | -0.170<br>(0.857)    | 3.971***<br>(0.992)  | 3.701<br>(2.438)     | 3.310***<br>(1.119)  |
| On Path                    | -0.490<br>(0.491)         | -0.879<br>(1.014)   | -0.022<br>(0.714)    | -0.423<br>(0.787)    | 1.544<br>(1.178)     | 1.190<br>(1.853)     | -1.862<br>(3.039)    |
| Off Path x Diff Max Payoff | 0.130<br>(0.264)          | 2.195<br>(1.722)    | 3.381***<br>(1.144)  | 5.743***<br>(0.815)  | 5.008***<br>(0.687)  | 3.489***<br>(1.130)  | -0.134<br>(1.855)    |
| On Path x Diff Max Payoff  | 0.000<br>(0.000)          | -0.298<br>(0.751)   | -1.438***<br>(0.457) | -2.563***<br>(0.459) | -4.285***<br>(0.699) | -1.930<br>(2.062)    | 2.060<br>(2.944)     |
| Intercept                  | 0.490<br>(0.491)          | 2.232***<br>(0.849) | 4.060***<br>(0.967)  | 9.883***<br>(1.145)  | 22.118***<br>(1.235) | 30.565***<br>(2.254) | 42.168***<br>(2.865) |
| N Participants             | 29                        | 23                  | 24                   | 60                   | 52                   | 24                   | 13                   |
| R <sup>2</sup>             | 0.00                      | 0.01                | 0.07                 | 0.06                 | 0.09                 | 0.04                 | 0.01                 |
| N                          | 1783                      | 1419                | 1476                 | 3682                 | 3198                 | 1474                 | 803                  |

(b)

Table 8: Beliefs Conditional on On- versus Off-Path Moves

*Notes:* This table shows how beliefs about suboptimal choices, expressed in percentage points, depend on whether the binary choice problem follows an on- or off-path move and on the difference in maximum continuation payoffs. Estimates are reported separately by participants' average reported probability of off-path moves: 0%, column (1); between 0% and 5%, column (2); between 5% and 10%, column (3); between 10% and 20%, column (4); between 20% and 30%, column (5); between 30% and 40%, column (6); and more than 40%, column (7). Subtable (a) estimates the specification corresponding to column (4) of Table 1 separately for each group of participants. Subtable (b) estimates the specification corresponding to column (5) of Table 1 separately for each group. The data comprise one-shot binary choice problems in part 1 and second-stage sequential choice problems in part 2 of the belief-elicitation experiment. Participants are grouped using their first-stage beliefs about off-path moves in part 2. Off Path and On Path are indicator variables equal to one if the binary choice problem follows an off- or on-path move, respectively, and zero otherwise. Diff Max Payoff denotes the difference in maximum continuation payoffs across the on- and off-path binary choice problems. Payoffs are measured in pounds. Standard errors clustered at the participant level are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.



**A.5. Response Times**

|                 | One-Shot<br>(1)     | Time<br>Sequential   |                     |                      |                     |
|-----------------|---------------------|----------------------|---------------------|----------------------|---------------------|
|                 |                     | (2)                  | (3)                 | (4)                  | (5)                 |
| Payoff Spread   | 0.231<br>(0.213)    | –                    | –                   | –                    | –                   |
| Payoff Level    | -0.082<br>(0.273)   | –                    | –                   | –                    | –                   |
| Diff Max Payoff | –                   | -0.543***<br>(0.155) | -0.327**<br>(0.148) | –                    | –                   |
| Diff Exp Payoff | –                   | –                    | –                   | -0.548***<br>(0.157) | -0.328**<br>(0.150) |
| Diff Min Payoff | –                   | –                    | -0.409**<br>(0.201) | –                    | -0.407**<br>(0.201) |
| Intercept       | 4.115***<br>(0.487) | 6.423***<br>(0.209)  | 6.370***<br>(0.199) | 6.423***<br>(0.210)  | 6.368***<br>(0.199) |
| R <sup>2</sup>  | 0.00                | 0.00                 | 0.00                | 0.00                 | 0.00                |
| N               | 2025                | 5912                 | 5912                | 5912                 | 5912                |

(a) Choices

|                 | One-Shot<br>(1)      | Time<br>Sequential   |                      |                      |                      |
|-----------------|----------------------|----------------------|----------------------|----------------------|----------------------|
|                 |                      | (2)                  | (3)                  | (4)                  | (5)                  |
| Payoff Spread   | -0.540<br>(0.674)    | –                    | –                    | –                    | –                    |
| Payoff Level    | 0.657<br>(0.881)     | –                    | –                    | –                    | –                    |
| Diff Max Payoff | –                    | -2.411***<br>(0.793) | -1.763**<br>(0.837)  | –                    | –                    |
| Diff Exp Payoff | –                    | –                    | –                    | -2.629***<br>(0.871) | -1.983**<br>(0.909)  |
| Diff Min Payoff | –                    | –                    | -1.220<br>(1.198)    | –                    | -1.048<br>(1.212)    |
| Intercept       | 10.422***<br>(1.423) | 26.310***<br>(1.280) | 26.147***<br>(1.240) | 26.367***<br>(1.306) | 26.186***<br>(1.252) |
| R <sup>2</sup>  | 0.00                 | 0.00                 | 0.00                 | 0.00                 | 0.00                 |
| N               | 2025                 | 5905                 | 5905                 | 5905                 | 5905                 |

(b) Beliefs

Table 9: Response Times and Payoffs; Levels

*Notes:* This table shows how response times in choices (subtable (a)) and belief reports (subtable (b)) vary with payoff features in one-shot binary choice problems, column (1), and in sequential choice problems, columns (2) to (5). Payoff Spread is the difference between the two payoffs in a binary choice problem. Payoff Level is the minimum payoff in the binary choice problem. Diff Max/Exp/Min Payoff denotes the difference in maximum/expected/minimum continuation payoffs between the on- and off-path binary choice problems in sequential choice problems. Payoffs are measured in pounds. Time is response time in seconds. The data comprise one-shot binary choice problems in part 1 and sequential choice problems in part 2 of the choice- and belief-elicitation experiments. Standard errors clustered at the participant level are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.

|                | Chose Suboptimally |                     | Chose Off-Path     |                     |
|----------------|--------------------|---------------------|--------------------|---------------------|
|                | Participant        | Participant         | Participant        | Participant         |
|                | Choice Freq.       | Avg. Belief         | Choice Freq.       | Avg. Belief         |
|                | (1)                | (2)                 | (3)                | (4)                 |
| Avg. Time      | 0.415<br>(0.317)   | 0.370***<br>(0.076) | 0.460<br>(0.299)   | 0.379***<br>(0.070) |
| Intercept      | 0.048<br>(1.647)   | 5.838***<br>(1.865) | 3.370**<br>(1.664) | 8.650***<br>(1.738) |
| R <sup>2</sup> | 0.02               | 0.14                | 0.02               | 0.18                |
| N              | 225                | 225                 | 225                | 225                 |

Table 10: Participants' Average Choices, Beliefs, and Response Times

*Notes:* This table shows how participants' average response time, measured in seconds, is related to their frequency of suboptimal choices in one-shot binary choice problems, column (1), and off-path moves, column (3), in the choice-elicitation experiment. It also shows how average response time is related to the average reported belief about others' suboptimal choices in one-shot binary choice problems, column (2), and off-path moves, column (4), in the belief-elicitation experiment. The data comprise one-shot binary choice problems in part 1 and sequential binary choice problems in part 2 of the choice- and belief-elicitation experiments. Heteroscedasticity-robust standard errors are shown in parentheses. \*, \*\*, and \*\*\* denote p-values < 0.1, < 0.05, and < 0.01, respectively.

## B. Instructions

Below, we reproduce screenshots of the instructions, practice rounds, paid rounds, and final questionnaire.

### B.1. Choice-Elicitation Experiment

#### Instructions Welcome!

Thank you for participating in this study.

In this study, we will ask you to make choices. You will get £2.55 when completing the study. You can earn an additional bonus payment of up to £2.55, depending on your choices.

This study is conducted under a strict policy of no deception. You will never be lied to during this study. The instructions will detail how the study works and how you are paid.

#### 'Bot'-Detection

This study is designed for humans and cannot be fulfilled using automated answers. You will be asked to transcribe easily legible words in this study. If you fail to transcribe a word three times, the study will be immediately terminated and you will not be able to complete it nor receive payment.

#### Informed Consent

Your participation in this study is voluntary. You can withdraw consent or discontinue participation at any time. All data and published work resulting from this study will be anonymised and maintain your individual privacy.

Please read this [information sheet](#) to decide if you want to participate in this study. You need to open this document to proceed. Should you have any questions or concerns, please contact the researchers or the UCL Research Ethics Office.

Participant's statement: I confirm that:

- a. I have read and understood the [information sheet](#);
- b. I understand that I can refuse or withdraw from this study at any time without affecting my rights or status;
- c. I am 18 or older and I freely consent to participate in this study;
- d. I may ask the researchers for a copy of these documents for my records.

Do you agree with the statement above and consent to participate in the study?

- ☐ Yes, I agree and wish to participate in the study.
- ☐ No, I do not agree and do not wish to participate in the study.

Next, we will provide instructions explaining how the study works and how you will be paid. As a reminder, you will never be lied to during this or any study at the UCL Experimental Laboratory for Finance and Economics (UCL ELFE).

When you are ready, please click "Next" to go on.

Next

# Instructions

In this study, you will make choices and you can earn an additional bonus payment of up to £2.55. Whether you earn that bonus payment will depend on your choices.

Click 'Next' to read the instructions.

Next

## Instructions

### Your Task

In each round you will face different sets of available balls. Each ball corresponds to a bonus amount.

Below, you see an example.

You can choose either the  ball or the  ball. If you choose the  ball, you receive a bonus payment of £0.65. If you choose the  ball, you receive a bonus payment of £1.15.



### Choices Between Boxes

In some rounds, you will first choose between two boxes.


After you choose a box, you can choose one of the two balls in that box.

**Important:** You can only choose a ball from the box you chose in the first step. You cannot choose a ball from the box that you did not choose.


Below, you see an example.

You can either choose Box 1, with the  and  balls, or Box 2, with the  and  balls.

Box 1



Box 2



### Multiple Rounds

In total, you will face 36 rounds. Your choices across rounds are independent and not related.

Each round will differ in the available balls and boxes.



At the end of the study, one round will be randomly selected, and your choice in that period will determine your bonus payment.

### Questions

You must answer the following questions correctly before you can proceed.

Box 1

Box 2



1. Which balls could you choose from if you were to choose Box 2 in the example above? Please select all that apply.
- ☐ Yellow and Green Ball.
  - ☒ Green and Orange Ball.
  - ☐ Orange and Purple Ball.
  - ☐ Purple and Yellow Ball.
2. If you had chosen the green ball and this round were selected for payment, which statement is true:
- ☒ You would earn a bonus of £2.05.
  - ☐ You would earn a bonus of £0.65.
  - ☐ You would earn a bonus of £2.15.
  - ☐ You would earn a bonus of £0.85.

Check Answers

All answers are correct! Click 'Next' to proceed.

Next

## Captcha Bot Detection

### Bot Detection - Attempt 1

Type the following word or phrase into the box below, then press 'Next'. Answers are not case-sensitive. You have three attempts. If you fail all three attempts, the task will end and you will not be paid. You have one minute per attempt.

sums

Next

## Practice Rounds

# Practice Round

You will now have 1 practice round to familiarise yourself with the study.  
You will receive feedback after your choices.  
You will not be paid for this round.

Next

# Practice Round

Box 1

£0.65

£1.25

Box 2

£0.75

£1.15

Please choose a box:

Box 1

Box 2

Next

Open Instructions

# Practice Round

Box 1

£0.65

£1.25

Box 2

£0.75

£1.15

Please choose a box:

Box 1

Box 2

Please choose a ball from Box 1:

Next

Open Instructions

48



## Practice Round Feedback



In the practice round, you chose Box 1 and the  ball.

If this were a paid round, you would be paid a bonus of £1.25.

Next

## Main Task Paid Rounds


You will now make choices in 36 rounds.

One of these rounds will be randomly selected for payment at the end of the study.

No feedback will be provided.

Next

## Round 1 out of 36



Please choose a ball:



Next

Open Instructions

Round 10 out of 36

Box 1

£1.45

£1.35

Box 1 ☐

Box 2

£1.55

£0.95

Box 2 ☐

Please choose a box:

Next

Open Instructions

Round 10 out of 36

Box 1

£1.45

£1.35

Box 1 ☐

Box 2

£1.55

£0.95

Box 2 ☒

Please choose a box:

Please choose a ball from Box 2:

Next

Open Instructions

Questionnaire

# Questionnaire

Before we conclude, we would like to ask you a few questions about yourself.

Please enter your age:

What is the highest level of education that you **completed**?

- ☐ Primary Education or Less
- ☐ Secondary Education/High School
- ☐ Business, technical, or vocational school after Secondary Education/High School
- ☐ Some college or university qualification, but not a Bachelor
- ☐ Bachelor or equivalent
- ☐ Master or Post-graduate training or professional schooling after college (e.g., law or medical school')
- ☐ Ph.D. or equivalent

Choose the field that best describes your **primary field of education**.

- ☐ Generic
- ☐ Arts and Humanities
- ☐ Social Sciences and Journalism
- ☐ Economics
- ☐ Education
- ☐ Business, Administration, and Law
- ☐ Computer Science, Information, and Communication Technologies
- ☐ Natural Sciences
- ☐ Mathematics and Statistics
- ☐ Engineering, Manufacturing, and Construction
- ☐ Agriculture, Forestry, Fisheries, and Veterinary
- ☐ Health and Welfare
- ☐ Services (Transport, Hygiene and Health, Security, and Other)

You must answer each question before you can continue.

Next

## B.2. Belief-Elicitation Experiment

### Instructions

#### Welcome

Thank you for participating in this study.

In this study, we will ask you to make choices. You will get £3.00 when completing the study.

You can earn an additional bonus payment of up to £3.00, depending on your choices.

This study is conducted under a strict policy of no deception. You will never be lied to during this study.

The instructions will detail how the study works and how you are paid.

#### 'Bot'-Detection

This study is designed for humans and cannot be fulfilled using automated answers. You will be asked to transcribe easily legible words in this study. If you fail to transcribe a word three times, the study will be immediately terminated and you will not be able to complete it nor receive payment.

#### Informed Consent

Your participation in this study is voluntary. You can withdraw consent or discontinue participation at any time. All data and published work resulting from this study will be anonymised and maintain your individual privacy.

Please read this [information sheet](#) to decide if you want to participate in this study. You need to open this document to proceed. Should you have any questions or concerns, please contact the researchers or the UCL Research Ethics Office.

Participant's statement: I confirm that:

- a. I have read and understood the [information sheet](#);
- b. I understand that I can refuse or withdraw from this study at any time without affecting my rights or status;
- c. I am 18 or older and I freely consent to participate in this study;
- d. I may ask the researchers for a copy of these documents for my records.

Do you agree with the statement above and consent to participate in the study?

- ☐ Yes, I agree and wish to participate in the study.
- ☐ No, I do not agree and do not wish to participate in the study.

Next, we will provide instructions explaining how the study works and how you will be paid. As a reminder, you will never be lied to during this or any study at the UCL Experimental Laboratory for Finance and Economics (UCL ELFE).

When you are ready, please click "Next" to go on.

Next

### Instructions

In this study, you will make choices and guesses about the decisions of other participants who have already completed this study.

You can earn an additional bonus payment of up to £3.00.

Whether you earn that bonus payment will depend on the accuracy of your guesses and your choices.

You will first learn how the other participants made their choices.

Keep in mind that you will not make these same choices but will guess the choices of others.

Click 'Next' to read the instructions of the other participants.

Next

# Instructions of the Other Participants

## Your Task

In each round you will face different sets of available balls. Each ball corresponds to a bonus amount.

Below, you see an example.

You can choose either the  ball or the  ball. If you choose the  ball, you receive a bonus payment of £0.65. If you choose the  ball, you receive a bonus payment of £1.15.



## Choices Between Boxes

In some rounds, you will first choose between two boxes.


After you choose a box, you can choose one of the two balls in that box.

**Important:** You can only choose a ball from the box you chose in the first step. You cannot choose a ball from the box that you did not choose.


Below, you see an example.

You can either choose Box 1, with the  and  balls, or Box 2, with the  and  balls.

Box 1



Box 2



## Multiple Rounds

In total, you will face 36 rounds. Your choices across rounds are independent and not related.

Each round will differ in the available balls and boxes.

At the end of the study, one round will be randomly selected, and your choice in that round will determine your bonus payment.

This is the end of the instructions of the other participants. Click on 'Next' to continue to a practice round of the other participants.

Next

Practice Round of the Other Participants' Task

Box 1

£0.65

£1.25

Box 2

£0.75

£1.15

Please choose a box:

Box 1

☐

Box 2

☐

Next

Open Instructions

Practice Round of the Other Participants' Task

Box 1

£0.65

£1.25

Box 2

£0.75

£1.15

Please choose a box:

Box 1

☒

Box 2

☐

Please choose a ball from Box 1:

☐


☐

Next


Open Instructions

## Practice Round of the Other Participants' Task

Box 1



Box 2



In the practice round, you chose Box 1 and the  ball.

If this were a paid round and the other participant had made this choice, they would be paid a bonus of £1.25.

Click 'Next' to see your instructions for this study.

Next

# Instructions

## Your Task

In this study, you will face 36 rounds. In each round, you will be **randomly matched with other participants that already completed the study**.

You will make guesses that will determine how likely you are to get a bonus payment of £3.00. You will see the same information as the other participants.

## Your Guesses

In each round, you are randomly matched with one of the other participants. You are then asked to guess the choice they made when they chose between boxes or balls.

The more accurate your guesses are, the higher is the probability of receiving the bonus payment of £3.00.

Below, you see an example.

Box



What do you think is the probability that the other participant chose:

%

%

The payment rule is designed so that you maximise your probability of winning the bonus payment if you tell us what you believe the probability is that a given ball or box was selected.

To ensure this, your payment from guesses is determined as follows:

- One guess will be randomly selected, all equally likely.
- Two numbers are randomly selected between 0 and 100, with all numbers being equally likely.
- You get a bonus of £3.00 if at least one of the two numbers is smaller than the probability you assign to the alternative chosen by the other participant you were randomly matched with.

This rule applies to guesses about both boxes and balls.

## Multiple Rounds

In total, you will face 36 rounds. Your guesses across rounds are independent and not related.

Each round will differ in the available balls and boxes.

At the end of the study, one round will be randomly selected, and one of your guesses in that round will determine your bonus payment. Each guess is equally likely selected.

## Questions



You must answer the following questions correctly before you can proceed.  
Questions 1 and 2 adapt the questions the other participants answered during the study.



1.  
Which balls could the other participant choose from if they had chosen Box 2 in the example above?
  - ☐ Yellow and Green Ball
  - ☐ Green and Orange Ball
  - ☐ Orange and Purple Ball
  - ☐ Purple and Yellow Ball
2.  
If the other participant had chosen the green ball and this round were selected for payment, which statement is true:
  - ☐ The other participant would earn a bonus of £2.05.
  - ☐ The other participant would earn a bonus of £0.65.
  - ☐ The other participant would earn a bonus of £2.15.
  - ☐ The other participant would earn a bonus of £0.85.
3.  
In this experiment, you are guessing what the other participants in this experiment chose. Choose the correct answer if you want to maximise the probability of winning the bonus payment of £3.00.
  - ☐ Your guess does not matter for earning the bonus.
  - ☐ It is always better to make the best guess possible.

Check Answers

## Captcha

### Bot Detection

#### Bot Detection - Attempt 1

Type the following word or phrase into the box below, then press 'Next'. Answers are not case-sensitive.  
You have three attempts. If you fail all three attempts, the task will end and you will not be paid.  
You have one minute per attempt.

sums

Next

Main Task  
Paid Rounds

You will now make choices in 36 rounds.  
One of these rounds will be randomly selected for payment at the end of the study.  
No feedback will be provided.

Next

Round 1 out of 36

Box

£0.55

£0.45

What do you think is the probability that the other participant chose:

%%

Next

Other Participants' Instructions

Your Instructions

Round 2 out of 36

Box

£1.35

£1.95

What do you think is the probability that the other participant chose:

0

%

100

%

Next

Other Participants' Instructions

Your Instructions

Round 10 out of 36

Box 1

£0.95

£1.05

Box 2

£0.45

£1.65

What do you think is the probability that the other participant chose:

Box 1

%

Box 2

%

Next

Other Participants' Instructions

Your Instructions

Round 10 out of 36

Box 1

£0.95

£1.05

Box 2

£0.45

£1.65

You predicted the following probabilities of the other participant choosing Box 1:

%

Suppose you've been matched with one of the other participants who chose Box 1.

What do you think is the probability that the other participant chose:

%

%

You predicted the following probabilities of the other participant choosing Box 2:

%

Suppose you've been matched with one of the other participants who chose Box 2.

What do you think is the probability that the other participant chose:

%

%

Next

Other Participants' Instructions

Your Instructions

## Questionnaire

## Questionnaire

Before we conclude, we would like to ask you a few questions about yourself.

Please enter your age:

What is the highest level of education that you **completed**?

- ☐ Primary Education or Less
- ☐ Secondary Education/High School
- ☐ Business, technical, or vocational school after Secondary Education/High School
- ☐ Some college or university qualification, but not a Bachelor
- ☐ Bachelor or equivalent
- ☐ Master or Post-graduate training or professional schooling after college (e.g., law or medical school')
- ☐ Ph.D. or equivalent

Choose the field that best describes your **primary field of education**.

- ☐ Generic
- ☐ Arts and Humanities
- ☐ Social Sciences and Journalism
- ☐ Economics
- ☐ Education
- ☐ Business, Administration, and Law
- ☐ Computer Science, Information, and Communication Technologies
- ☐ Natural Sciences
- ☐ Mathematics and Statistics
- ☐ Engineering, Manufacturing, and Construction
- ☐ Agriculture, Forestry, Fisheries, and Veterinary
- ☐ Health and Welfare
- ☐ Services (Transport, Hygiene and Health, Security, and Other)

You must answer each question before you can continue.

Next