

Contents lists available at [ScienceDirect](#)

Journal of Economic Behavior and Organization

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

Take it or leave it: Experimental evidence on the effect of time-limited offers on consumer behaviour

Robert Sugden^{a,*}, Mengjie Wang^a, Daniel John Zizzo^b

^a Centre for Behavioural and Experimental Social Science, Centre for Competition Policy, and School of Economics, University of East Anglia, Norwich NR4 7TJ, United Kingdom

^b School of Economics, University of Queensland, 39 Blair Drive, St Lucia QLD 4067, Australia

ARTICLE INFO

Article history:

Received 18 October 2017

Revised 9 July 2019

Accepted 9 September 2019

Available online xxx

JEL classifications:

C91 (design of experiments: laboratory, individual behaviour)

D83 (information, knowledge and uncertainty: search)

D91 (micro-based behavioural economics: role and effects of psychological, emotional, social and cognitive factors on decision-making)

Keywords:

Time-limited offers

Time constraints

Consumer search

Consumer choice

ABSTRACT

Making time-limited offers is a common retail pricing strategy. Economic theory implies that such offers inhibit price search, making markets less competitive. We investigate experimentally whether this effect is intensified by behavioural factors – specifically, feedback-conditional regret, reduced decision quality due to time constraints, and aversion to small-scale risk. Participants choose from a sequence of alternative price offers, one of which might be time-limited, under various conditions. These price search problems were matched with equivalent, time-unconstrained binary choices between lotteries. We find no evidence of regret effects. Surprisingly, time-limited offers are more likely to be chosen when the time available for decision-making is longer. Overall, individuals show aversion to small-scale risk; this is stronger in price search than lottery choice. Allowing for this, choices in the two types of task tend to be mutually consistent at the individual level, even when decision-making is subject to tight time constraints.

© 2019 The Author(s). Published by Elsevier B.V.
This is an open access article under the CC BY license.
(<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

A common tactic by salespeople is to make offers that (it is claimed) will be withdrawn unless accepted immediately. Familiar examples include the doorstep seller who claims that he is currently ‘in the area’ but will not be returning, the telephone seller who makes a ‘special offer’ that can be accepted only during that phone call, the internet site which offers a buy-now discount, and the seller of a used car who claims that another buyer has shown great interest in it and will be returning shortly. Writers who have infiltrated businesses report that sales staff are routinely instructed to use such tactics (Cialdini, 2001, p. 208; Bone, 2006, pp. 71–73). There is growing concern that, by using cookies to track the identities of website visitors, internet sellers may remove low-price offers between a potential buyer’s first and second visit. Viewed in

* Corresponding author.

E-mail address: r.sugden@uea.ac.uk (R. Sugden).

<https://doi.org/10.1016/j.jebo.2019.09.008>

0167-2681/© 2019 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license.

(<http://creativecommons.org/licenses/by/4.0/>)

Please cite this article as: R. Sugden, M. Wang and D.J. Zizzo, Take it or leave it: Experimental evidence on the effect of time-limited offers on consumer behaviour, Journal of Economic Behavior and Organization, <https://doi.org/10.1016/j.jebo.2019.09.008>

the perspective of competition regulation, such *time-limited* (or *exploding*) offers create barriers to the search processes by which consumers compare prices.¹

It has long been known that the presence of search costs makes markets less competitive and induces higher prices (Diamond, 1971; Salop and Stiglitz, 1977). Armstrong and Zhou (2016) present a range of theoretical models in which opportunities to make time-limited offers have similarly anti-competitive effects. In these models, profit-maximising monopolistic or oligopolistic firms use time-limited offers to deter search, with the result that there is an increase in prices and/or a reduction in the degree of match between consumers' preferences and the products they buy.² For a time-limited offer to have these *search-detering* effects, consumers' opportunities to investigate alternatives during the offer period must be significantly constrained, which will typically be the case only if that period is quite short and is determined separately for each consumer.³ In Armstrong and Zhou's models, consumers search offers sequentially, and so no other offers can be investigated during the period in which a time-limited offer is active. In this paper, we will be concerned with time-limited offers of this kind.

When offers are searched sequentially, a time-limited offer (if better than all preceding offers that are still open) presents a consumer with a binary decision problem under uncertainty. Accepting the offer gives her the certainty of a specific product at a specific price. Rejecting the offer gives her an uncertain prospect – the distribution of possible outcomes resulting from a continuation of the search process. Any search-detering properties of time-limited offers must be mediated by consumers' responses to such binary decision problems. Holding constant the behaviour of other firms, any given firm has a stronger incentive to make its offers time-limited, the greater the tendency for consumers to prefer certainty to uncertainty when engaging in search problems. In Armstrong and Zhou's models, consumers are risk-neutral agents who choose search strategies that maximise expected utility, given knowledge of the distributions from which offers are drawn. As Armstrong and Zhou (2016, p. 50) point out, it would be useful to investigate whether the search-detering effects that their models describe are intensified or mitigated by properties of consumers' choice behaviour that have been identified by behavioural research but which do not feature in traditional decision theory. We report an experiment that was designed to investigate behavioural mechanisms that might impact on consumers' choices between accepting and rejecting time-limited offers, and that might be expected to favour the choice of time-limited offers.

Under plausible assumptions about utility functions and if the value of the relevant purchase is low relative to a consumer's wealth, models of risk-neutral search closely approximate the implications of expected utility theory. However, there are a number of behavioural reasons for expecting consumers' responses to time-limited offers to be significantly risk-averse, favouring the choice of such offers. First, a consumer who rejects a time-limited offer and continues to search may find that the rejected offer was in fact the best available, and this may induce painful feelings of regret. Choosing the offer and terminating the search process can be a method of avoiding regret. Second, time-limited offers are often presented in ways that give consumers only a short time in which to decide between acceptance and rejection. In response to such time constraints, consumers may use simplifying heuristics that favour certainties. Third, experimental evidence often reveals high degrees of risk aversion in decisions that involve very small stakes. Our experiment investigates the effects of these three mechanisms on consumers' responses to time-limited offers.

We have deliberately chosen to focus on mechanisms that operate through behavioural properties of individual (rather than strategic or interactive) decision-making by consumers. We recognise that, in natural-world markets, consumers' propensities to accept time-limited offers might also be affected by judgements about the fairness of sellers' intentions, but that is outside the scope of our investigation. Our experimental design isolates the effects of non-interactive decision-making by using human subjects only as consumers.

All offers in our experiment, whether free-recall or time-limited, were generated exogenously. Given the research questions we want to address, this feature is important both for experimental control and for external validity. Since our interest is in how consumers respond to sellers' offers, we wanted to have as much control as possible over our participants' beliefs about the probabilities of the different types and values of offer they might face. We implemented this control by using random processes with fixed probability distributions, by explicitly informing participants about these distributions, and by allowing participants to learn about the distributions through repeated experience. Of course, a complete theory about retail markets would need to combine an explanation of how consumers respond to given offers with an explanation of how sellers choose what offers to make, given how consumers can be expected to respond. But we believe that it is productive to investigate these components separately, particularly because the experimental method is more suitable for investigating the behaviour of consumers than that of firms. Every experimental participant can be expected to be familiar with problems of price search on computer interfaces, but the external validity of decisions made by participants acting as firms is more

¹ Evidence about these sales strategies is reviewed by Office of Fair Trading (2010) and Armstrong and Zhou (2016). Armstrong and Zhou refer to current controversies about the alleged misuse of cookies by websites selling airline tickets, but take no view about the truth of these allegations – perhaps because it is not clear that sellers would benefit by imposing time limits that are not announced in advance.

² A related literature investigates time-limited offers in matching markets, such as entry-level labour markets for lawyers, physicians and academics. It has been shown both theoretically (Roth and Xing, 1994) and experimentally (Niederle and Roth, 2009) that if time-limited offers are permitted, matching markets tend to 'unravel' – that is, contracts are made increasingly early and the quality of matching declines.

³ A contrasting type of time-limited offer is exemplified by special promotions in supermarkets and by 'sales' periods in department stores. Such offers, which are made available to all consumers for periods of several days or even weeks, may be components of a strategy of varying prices over time in order to discriminate between consumers with different propensities to search (Varian, 1980), to manipulate consumers' reference points (Heidhues and Kőszegi, 2014), or to counteract consumers' tendencies to procrastination (O'Donoghue and Rabin, 2001).

questionable. One might also have concerns about how far, when experimental participants act as buyers, their attitudes to the fairness of offers made by fellow-participants are predictive of their attitudes to offers made by commercial retailers.⁴

In the rest of the paper, we begin with a brief review of previous research (Section 2). We then explain the behavioural mechanisms we will investigate, outlining the basic principles of our experimental design and stating the main hypotheses that we will test (Section 3). We go on to describe the implementation of our experiment in more detail (Section 4) and to report our initial results (Section 5). As our findings about the effects of time constraints were unexpected, we ran a second experiment to investigate time constraints in more detail (Section 6). To anticipate our conclusions (Sections 7 and 8), we find no evidence that the tendency for consumers to choose time-limited offers is intensified by desires to avoid regret. Surprisingly, tighter time constraints made time-limited offers less likely to be chosen. However, our subjects' behaviour was predominantly risk-averse. Overall, choices between accepting or rejecting time-limited offers were more risk-averse than when the same decision problems were framed as choices between lotteries.

In so far as our experimental results have external validity, the implication is that models of consumer behaviour that assume the validity of expected utility theory understate the attractive force exerted by time-limited offers. It is natural to conjecture that, when used in theories of retail competition, this assumption also understates the anti-competitive effects of such offers. However, our focus is on empirical properties of consumer behaviour.

2. Previous research

A number of previous experiments have investigated the effects of time-limited offers on buyers' search behaviour.

Huck and Wallace (2015) report an experiment in which subjects played the role of consumers buying a homogenous product sold by two (computerised) 'shops' at prices that were independently drawn from a common distribution. Each visit to a shop incurred a cost. In the baseline condition, each shop quoted a simple per-unit price, which was revealed only when the shop was visited. These were *free-recall* offers – that is, a consumer who returned to a previously-visited shop would find the original price (and no other) still available. Five other conditions represented different and more complex 'price frames'. In the 'time-limited offer' frame, the offer at the first shop visited had to be accepted or rejected at that visit; the second shop's offer was made with free recall. If the consumer returned to the first shop, a different price would be generated (also with free recall). The experiment found over-search in the baseline condition (i.e. the frequency with which the first offer was accepted was lower than would have been the case for rational risk-neutral agents) but under-search in the time-limited offer frame. This observation suggests that time-limited offers may be more attractive to real consumers than to their rational counterparts. However, Huck and Wallace (pp. 32–33) note that many subjects behaved as if they did not realise that the first shop's second offer could be lower than its initial time-limited offer. This misunderstanding may have induced a bias towards the choice of time-limited offers.

Brown, Viriyavipart and Wang (2018) report an experiment that implemented a variant of the Armstrong and Zhou (2016) duopoly model, but using incentivised human subjects to represent firms. Firms could choose whether or not to make their offers time-limited or free-recall. In one treatment, consumers were represented by human subjects; in another, they were computer programs that implemented optimal (risk-neutral) search strategies. The experiment found over-searching: at the first firm visited, time-limited offers were chosen less frequently by human subjects than by optimal search programs. Brown et al. (2018, p. 183) discuss two possible explanations of this result. One is that human buyers perceive time-limited offers as unfair, and react to this unfairness by rejecting such offers even when accepting them would be beneficial. The other is that, for cognitively constrained buyers who are forced to make decisions under tight time constraints, the 'fast, intuitive response to an exploding offer is to run from the offer and search further'. Brown et al. favour the second explanation, on the ground that rejection of beneficial time-limited offers (i.e. offers that a risk-neutral rational buyer would accept) was negatively correlated with subjects' scores on a Cognitive Reflection Test, and that in exit surveys, subjects who had been buyers did not express concerns about the fairness of time-limited offers. This experiment also found that acceptance of non-beneficial time-limited offers was positively correlated with subjects' levels of risk aversion, as elicited in a post-experiment questionnaire.

Tang, Bearden and Tsetlin (2009) report an experimental game between a 'proposer' (corresponding with a firm in the context of search by consumers) and a 'responder' (consumer). In this game, the responder sees a sequence of randomly-generated free-recall offers. Before this sequence begins, the proposer makes an exogenously fixed offer that is intermediate between the best and worst offers that can appear in the random sequence. The responder must choose one of the offers. The proposer chooses a 'deadline' for her offer – that is, the stage in the random sequence at which her offer must be either accepted or rejected. Her objective is to maximise the probability that her offer will be accepted. The experiment was primarily designed to investigate the behaviour of proposers relative to a benchmark of rational play, but it also produced evidence about responders' willingness to accept time-limited offers. Median responses were close to the rational-play benchmark, but there was a large amount of unexplained variation, perhaps reflecting the fact that, since each participant played the game only once, there were no opportunities for experiential learning. Lau, Bearden and Tsetlin (2011) report

⁴ A further consideration is that, for the kind of price search situations we consider, the seller's optimisation problem is much more complex than the consumer's. See, for example, Lippman and Mamer (2019), whose theoretical model is quite similar to our experimental set-up, except that one offer (the counterpart of our 'red card' – see later) is made by a 'strategic buyer' who chooses whether it is free-recall or time-limited. Lippman and Mamer conclude: 'Our analysis and examples reveal that the strategic buyer's choice of offer duration is complex: almost anything can occur' (Abstract).

an experiment with a similar design, but in which the proposer makes a binary choice between imposing or not imposing a deadline, and the responder is able to punish a proposer after accepting her offer. There was a strong tendency for punishment to be directed at proposers who had set deadlines.

These previous experiments provide mixed evidence about the relative attractiveness of time-limited and free-recall offers in environments in which a range of behavioural mechanisms may be at work simultaneously. Our experiment was designed to find out more about consumers' responses to time-limited offers by isolating the specific effects of a restricted set of mechanisms.

3. Basic principles of experimental design and hypotheses to be tested

Our experiment had two parts. In Part 1, each subject faced thirty *price search* tasks, of which fifteen (randomly positioned in the series of tasks, and not announced in advance) involved time-limited offers. In Part 2, each subject faced fifteen *lottery* tasks requiring binary choices between lotteries. In the price search tasks, the price offers between which subjects had to choose were random draws, made independently for each subject, for each task and for each offer, from pre-specified distributions which remained constant across the thirty tasks. Lottery tasks were constructed separately for each subject, to match the specific distributions of offers that subject had faced in the fifteen price search tasks with time-limited offers. Viewed in the framework of expected utility theory, each lottery task was equivalent to the corresponding price search task. Subjects were not told about this correspondence and, because the frames were so different, it would have been very difficult to detect.⁵ To allow between-subject tests of the effects of regret, each experimental session was randomly assigned to one of two treatments, which differed in terms of the feedback provided after a time-limited offer had been accepted in a price search task. To allow within-subject tests of the effects of time constraints, both treatments incorporated variation in the time allowed for making a decision about a time-limited offer.

The price search tasks were designed to provide a controlled and stylised representation of environments in which consumers engage in price search and in which time-limited offers appear relatively infrequently and without prior notice. In each task, the subject was given a 'budget' and instructed to buy a 'good' by spending from this budget; any unspent surplus constituted her earnings from the task. The subject was able to see six price offers, which appeared sequentially on a computer screen with short time intervals between them. In most cases, offers were free-recall. Although free-recall offers could be accepted at any time, the subject was free to wait until all such offers had appeared and then choose whichever of these she judged best (presumably the lowest). In fifteen of the tasks, however, one of the first three offers to appear would be flagged as time-limited. In this case, the subject was able to accept the offer only in the interval before the next offer appeared.

We gave subjects full information about the random processes that generated the six price offers in each task. These processes were given an intuitive interpretation in terms of 'deals' or 'cards' (see later) and were held constant throughout the experiment. Thus, the actual values of the offers in any task provided no information about offers in later tasks. In this sense, repetition did not provide any opportunities for learning, other than in terms of general familiarity with the task. However, although we expected that subjects would understand how the offers were generated, we did not expect that they would be able to calculate optimal solutions to the problems they faced: such calculations would be mathematically challenging even in the absence of the time constraints our design imposed. Our design strategy was to give each Part 1 task the 'feel' of a natural price search problem rather than to mimic the properties of theoretical models of search. By framing the structure of the task in simple terms and repeating it many times, we tried to ensure that subjects would gain an intuitive understanding of it and converge on patterns of response similar to those that they would use in natural price search problems. We take it that the theories of search that we test are intended to predict the behaviour of ordinary consumers in everyday settings.⁶

The experiment was designed to investigate three behavioural mechanisms that might be expected to make real consumers more likely than their counterparts in models of rational risk-neutral search to choose time-limited offers.

3.1. Feedback-conditional regret

The hypothesis that decision-making behaviour is influenced by anticipated regret was proposed by Bell (1982) and Loomes and Sugden (1982). Cues which prompt individuals to anticipate possible future regret have been found to make individuals less likely to take risks in consumer choice (Simonson, 1992) and in sexual behaviour (Richard et al., 1996). In the context of time-limited offers, the dependence of regret on feedback about the outcomes of non-chosen options, theorised by Humphrey (2004), is particularly significant.

Consider a choice between two lotteries with monetary outcomes, defined on the mutually exclusive and exhaustive events E_1, \dots, E_n . The 'safe' lottery S pays s in every event; the 'risky' lottery R pays r_1, \dots, r_n , depending on which event

⁵ The random draws of offer prices were represented to subjects as 'deals' of cards, made at the start of each task. That representation reflected the actual sequence of operations by the computer program: offer prices for each task were determined only as that task appeared on a subject's screen. Given this feature of the experiment, price search tasks necessarily preceded lottery tasks. As we explain later, it seems unlikely that our main results are affected by the order of the two types of task.

⁶ This general methodological strategy is explained and defended by Sitzia and Sugden (2011).

obtains. First suppose that, whichever lottery is chosen, the true state will be revealed. Then, according to regret theory, when the decision-making agent considers choosing S , she will anticipate feelings of regret in those events E_i where $s_i < r_i$; these anticipations have a negative impact on the expected utility of S . Conversely, when she considers choosing R , she will anticipate regret in those events E_j where $r_j < s_j$; these anticipations have a negative impact on the expected utility of R .⁷ Now suppose instead that the true state will be revealed only if R is chosen. In this case, as before, choosing R exposes the agent to the possibility of regret-inducing feedback; but now choosing S cuts off such feedback. Thus, a regret-averse agent might choose S if there will be feedback but R if there will not be. Zeelenberg et al. (1996) find evidence of this pattern of behaviour.

The Dutch Postcode Lottery provides an extreme example of how this tendency can be exploited by sellers. By drawing a winning postcode rather than a winning ticket number, this lottery design exposes individuals who do not bet to the possibility of extreme regret if their neighbours bet and win. Zeelenberg and Pieters (2004) find that anticipated regret is positively correlated with preferences for participating in a postcode lottery rather than a conventional alternative. Time-limited offers may exploit a similar tendency. Rejecting a time-limited offer exposes a consumer to the possibility of regret, but if offers are discovered only by active search, a decision to stop searching cuts off feedback that could cause regret about having accepted such an offer.⁸

To investigate this possibility, we allocated each subject to one of two different feedback treatments. In the *No Feedback* treatment, if an offer was accepted before all six had been shown, the remaining offers were not revealed. Thus, in accepting a time-limited offer, a subject cut herself off from feedback that could reveal that she would have done better if she had rejected it. In the *Regret Feedback* treatment, whenever an offer was accepted, any remaining offers were immediately displayed. We tested the following hypothesis:

Hypothesis 1. Other things being equal, the probability with which time-limited offers are chosen is greater when feedback is absent than when it is present.

3.2. Time constraints

Consumers' choices between accepting and rejecting time-limited offers may be affected by the *time constraint* of having to make a decision in a short period of time, rather than being able to think about the decision more carefully. Although time constraints and search deterrence often go together, they are conceptually distinct mechanisms. In some markets (for example, for houses in particular locations), the interval between the arrival of offers may be long enough for buyers to make fully considered decisions amongst all currently available offers, but sellers may still be able to use time-limited offers to deter buyers from waiting for further offers to arrive before making a decision. Conversely, consider a market in which the products supplied by incumbent firms and the prices at which these can be bought are well-known. A door-to-door salesperson offering a new product might still use time-limited offers to impose a time constraint. Our experiment is designed to distinguish between these two mechanisms.

It is also important to distinguish between objective time constraints and subjective 'time pressure' – the experience of anxiety that can be induced by such constraints. Some familiar forms of time-limited offers are presented in ways that are intended to create or amplify time pressure. In internet selling, it is common for sellers to induce time pressure by using messages about the limited number of offers available and about other buyers' current interest in them. When time-limited offers are made in face-to-face (or voice-to-voice) interactions, such as doorstep and telephone selling, conversational norms add to the pressure for quick responses. Our experiment focuses on the objective dimension of time constraints.⁹

The psychological literature suggests that people often respond to time constraints by using relatively simple heuristics that are adapted to maintaining decision accuracy with reduced cognitive effort, rather than by using truncated forms of reasoning processes that would generate correct decisions in the absence of these constraints. Payne et al. (1988) find that, under severe time constraints, experimental subjects making choices under risk focus on a subset of the available information and change their information-processing strategies. Finucane et al. (2000) find that time constraints increase subjects' reliance on the 'affect heuristic', which induces a negative correlation between judgements of risk and benefit (for example, a tendency to underestimate the radiation risk from the use of x-rays in hospitals).

There is some experimental evidence that, when making choices under risk, people respond to time constraints by using modes of reasoning that tend to accentuate risk aversion. In an investigation of choices between pairs of gambles under different levels of time constraint, Ben Zur and Breznitz (1981) find that subjects spend proportionately more time attending to negative dimensions of gambles under tight time constraints and so make more risk-averse decisions. Kocher et al. (2013) find that, although time constraints have no effect on risk attitudes revealed in choices between prospects offering gains, they increase risk aversion in choices between prospects which may impose losses. In a costly

⁷ For simplicity, we ignore anticipated feelings of 'rejoicing' (the opposite of regret). In the regret literature it is common to assume that regret is a much stronger emotion than rejoicing.

⁸ It is possible that decisions are influenced by anticipated regret even when no relevant feedback will be provided. (Hearing that *someone* has won a lottery might be sufficient to induce regret about not buying a ticket.) The effects of such regret would be the same in both our feedback treatments and in both price search and lottery tasks and would be picked up in our measures of risk aversion. Our interest is in the specific effects of feedback.

⁹ Future research could look at the effect of subjective 'time pressure', for example by using relevant psychological scales.

search task with free recall, [Ibanez et al. \(2009\)](#) find that participants tend to search less when they have little time to choose about whether to terminate search, but that this effect dissipates as participants gain experience of the task. In an investigation of lottery choice under extreme time constraints, [Nursimulu and Bossaerts \(2014\)](#) find that reducing the time available for decision-making from 5 s to 1 s increases risk aversion, as modelled in prospect theory by the curvature of the utility function for gains, but also has the offsetting effect of reducing the distortion of probability information. However, [Saqib and Chan \(2015\)](#) report a lottery choice experiment in which time constraints induce risk aversion for losses but risk-seeking for gains; their explanation for this pattern is that, under time constraints, subjects focus on the most extreme positive and negative outcomes. Overall, the evidence is not clear-cut, but it suggests a psychological mechanism by which time constraints might increase consumers' willingness to accept time-limited offers.

We investigated the effect of time constraints by randomising the time interval between the appearance of offers, and hence the length of time for which time-limited offers were available. The time interval (which was constant within a task) could be either 4 s or 12 s. Given the simplicity and repetitive nature of the tasks, we expected that 12 s would be sufficient for most subjects to make a considered (although not necessarily optimal) decision, but that 4 s would be perceived as a significant time constraint.¹⁰ We will refer to 4 s as a *tight* constraint, and 12 s as a *loose* one. We tested the following hypothesis:

Hypothesis 2. Other things being equal, the probability with which time-limited offers are chosen is greater when time constraints are tighter.

3.3. Risk aversion

In experiments involving choice amongst lotteries with low-value consequences, subjects often reveal high degrees of risk aversion. [Rabin \(2000\)](#) shows that, if conventional expected-utility models are fitted to the patterns of behaviour typically observed in such experiments, these models imply manifestly unrealistic degrees of risk aversion for decisions involving higher (but still modest) stakes. This inconsistency can be eliminated if, as in prospect theory, utility is defined as a function of changes in wealth relative to a reference point, rather than as a function of wealth itself, since this allows there to be particularly high degrees of risk aversion in the neighbourhood of any reference point ([Kahneman and Tversky, 1979](#)). If individuals are significantly averse to small risks, time-limited offers may be more attractive to consumers than is implied by models of risk-neutral search. Hence the following hypothesis:

Hypothesis 3. In choosing whether or not to accept time-limited offers, experimental subjects tend to reveal risk aversion.

Given the degree of risk aversion typically observed in experiments, confirmation of Hypothesis 3 might be thought unsurprising. There is a good deal of existing evidence that in search tasks, both with and without recall, experimental subjects tend to search less than would be optimal for risk-neutral agents (e.g. [Rapoport and Tversky, 1970](#); [Sonnemans, 1998](#); [Seale and Rapoport, 1997](#); [Schunk and Winter, 2009](#)).¹¹ It is perhaps more interesting to ask whether individuals display the same attitudes to risk when deciding whether to reject a time-limited offer as they do when choosing whether to enter a risky lottery. The equivalence between these two choice problems might be less salient to most people than it is to decision theorists. Thus, if individuals use simplifying heuristics, the heuristics primed by price search problems might not be the same as those primed when decision problems are explicitly framed as choices between lotteries. The lottery tasks allow us to investigate this issue. Since behavioural theory does not provide unambiguous predictions about the direction that differences between the two tasks might take, we tested the following directionless hypothesis:

Hypothesis 4. Risk attitudes revealed in decisions about accepting or rejecting time-limited offers are systematically different from those revealed in binary choices between certainties and lotteries.

4. Design details and implementation

4.1. Overall structure of experiment

As explained above, each session of the experiment was randomly assigned to one of the two treatments (No Feedback or Regret Feedback). Part 1 of each session consisted of thirty price search tasks; Part 2 consisted of fifteen lottery tasks. At the beginning of each part, each subject received a copy of the instructions for that part; these instructions were read aloud by the experimenter. These instructions are reproduced in Appendix 1. Each subject then completed a computerised questionnaire which tested her understanding of the tasks. If a subject made a mistake, the computer would show her the correct answer and the relevant part of the instructions. Subjects were invited to ask the experimenter for clarification.

In the price search tasks, subjects' budgets and offer prices were expressed in an experimental currency unit, *experimental points* (EP). Each subject's earnings from each task (conditional on the task being selected for payment: see later) were

¹⁰ In the context of doorstep, telephone or internet selling, a requirement to make a decision within 12 s might be perceived as an extreme time constraint. But the decision problems typically faced in such contexts are much less simple and repetitive than those in our experiment.

¹¹ Schunk and Winter find no significant relationship between search behaviour and a measure of risk attitude. They suggest that observed under-searching may be due to loss aversion rather than risk aversion.

equal to the budget (always 10EP) minus the price of the accepted offer (never greater than 10EP). Thus, the subject had an incentive to choose the offer with the lowest price. Each lottery task was a binary choice between two lotteries with outcomes expressed in experimental points. 'Lottery 1' gave a stated outcome for sure. 'Lottery 2' had five possible outcomes, each with a probability of 0.2.

At the end of each session, the computer randomly selected two price search tasks and one lottery task for each subject. Each subject's earnings were the sum of her earnings from the two selected price search tasks and the outcome of the lottery she had chosen in the selected lottery task. If she had chosen Lottery 2, its outcome was determined by a random process, using the stated probabilities. Subjects were paid at the exchange rate of £1 for every 2.5EP. This payment mechanism was used because it satisfied two design requirements. Subjects did not receive instructions for Part 2 until Part 1 had been completed, but we wanted them to perform the Part 1 tasks with full information about the incentives for these tasks. Since we needed to compare responses to matched pairs of Part 1 and Part 2 tasks, we wanted the incentives to be the same for every task in the experiment, but there were twice as many tasks in Part 1 than in Part 2.

4.2. Price search tasks

The basic structure of these tasks was described in Section 3 above; here we fill in the details. Offers were represented on subjects' screens by pictures of 'cards'. As with a pack of playing cards, all cards looked identical when viewed face down. When viewed face up, each card showed a price offer. At the start of a task, a subject saw a row of six cards, all face down. In terms of the prices they might offer, there were two types of card, *red* and *blue*. In all tasks, one of the cards was red and five were blue. The red card had been randomly assigned to one of the first three positions in the row, but the subject was not told which position this was. The offer on each blue card was a price drawn independently and randomly from the set {0.00EP, 0.01EP, ..., 10.00EP}. The offer on the red card was a price drawn at random from the set {0.00EP, 0.01EP, ..., 4.00EP}. Thus, the red card was particularly likely to show a low-price offer. Subjects were fully informed about the distributions of offers on the two types of card. Our reasons for using these particular parameters will be explained later.

During each task, the cards were turned over by the computer, one by one and working from left to right. The time interval between cards being turned over was fixed within each task at either 4 s or 12 s. Which interval was used in each task was randomised, independently for each participant, subject to the constraint that each participant faced each interval in fifteen tasks. There was no time limit for completing the task as a whole. Subjects could not move on to Part 2 of the experiment until everyone in the session had completed Part 1. Thus, an individual subject who chose not to see all six offers was unlikely to save time by doing this.

The offers on blue cards were always free-recall. Whenever a free-recall card (of either colour) was turned over to display its offer, a button appeared below it on the screen, labelled 'Click to choose this offer', together with the message 'This offer will stay available throughout the task'. The card remained turned over until the end of the task. The subject was free to choose that offer at any time from then on, by clicking on the button. Thus, it was possible for a subject to wait until all cards had been turned over and then to choose from the complete set of free-recall offers, with no time constraint.

The offer on the red card could be either free-recall (described to subjects as a 'standard red card') or time-limited (a 'time-limited red card'). We will classify tasks as 'free-recall' or 'time-limited' according to the properties of the red card. Whether the red card offer was free-recall or time-limited was randomised for each task, independently for each participant, subject to two constraints: that each subject faced fifteen free-recall tasks and fifteen time-limited tasks; and that in each of these sets of fifteen tasks, the split between 4 s and 12 s time intervals was either 7:8 or 8:7. Subjects were not told which type of offer the red card would make until it was turned over. If the offer was time-limited, a 'Click to choose this offer' button appeared below the card, together with the message 'This offer will stay available for 4 seconds' or 'This offer will stay available for 12 seconds'. A countdown clock showed the number of seconds remaining. At the end of this time interval, if the offer had not been accepted, the card was turned back, making the offer price no longer visible, and its choice button, availability message and countdown clock disappeared. Simultaneously, the next card in the sequence was turned over.

In price search tasks, the only difference between the two treatments occurred after an offer was accepted. In the No Feedback treatment, the offer price that the subject had chosen was shown on the screen, and the task ended at that point. Thus, if the subject chose an offer before all the cards had been turned over, she could not get any information about the offers on the remaining cards. In the Regret Feedback treatment, any remaining cards were turned over, revealing their prices. The subject was not allowed to change her decision at this stage.

Fig. 1 shows a typical screenshot for a time-limited task. In this example, the third card has just been turned over. It is a time-limited offer with a price of 1.75EP. The subject has 3 s remaining in which to choose whether to accept this offer. The two preceding free-recall prices are 2.27EP and 4.09EP.

Viewed in the framework of standard decision theory, free-recall price search tasks are trivial, irrespective of a subject's attitude to risk. In such a task, a rational subject would wait until all prices had been revealed and then choose the lowest. A time-limited task is also trivial if the time-limited price is greater than or equal to any preceding free-recall offer: in such a case, the time-limited offer would always be rejected. A significant decision problem occurs only when a time-limited offer is better than any preceding offer and (as was always the case in our design) other offers remain to be revealed. In this case, there are only two relevant options for a rational subject: either to accept the time-limited offer; or to reject it,

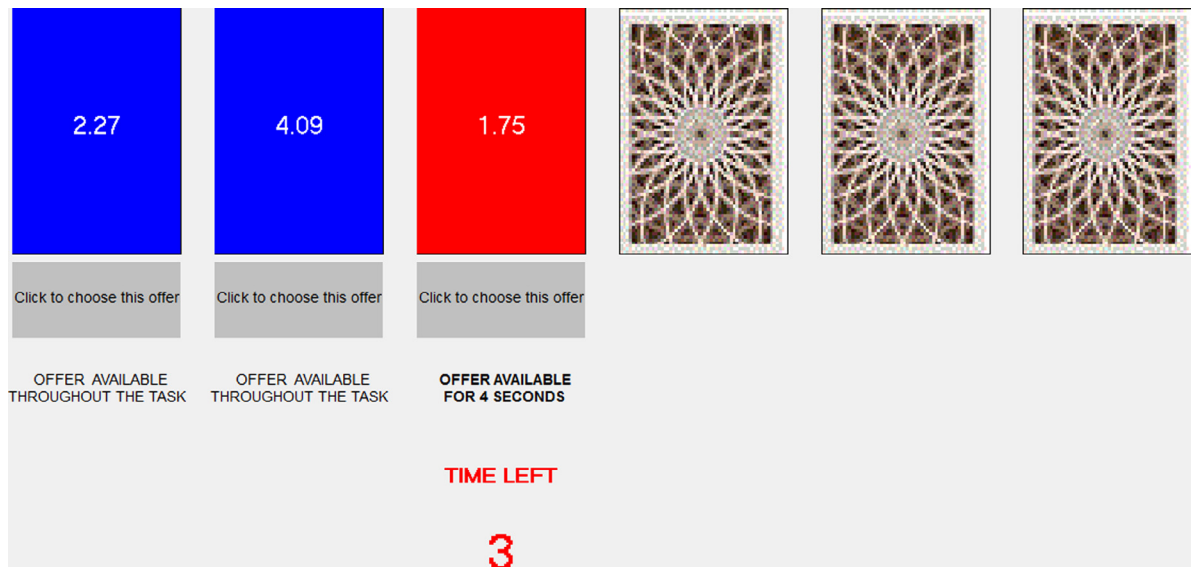


Fig. 1. Screen shot of a typical time-limited task.

wait until all offers have been revealed, and then choose the best free-recall offer. Time-limited tasks which present such problems will be called *consequential*.

For reasons of statistical power, we needed to use parameters that would ensure that, in the experiment as a whole, approximately equal numbers of subjects would accept and reject time-limited offers in consequential tasks. We therefore chose the parameters so that, conditional on facing a consequential task, there was a 0.5 probability that a risk-neutral expected-utility-maximising subject would accept the time-limited offer. For similar reasons, we needed that, in a significant proportion of consequential tasks, the expected values of 'accept' and 'reject' would be relatively close to one another, so that subjects could be expected to find the decision problem relatively difficult. In order to achieve this objective, it was necessary that time-limited offers appeared early in the offer sequence and, on average, had lower prices than free-recall offers. The fact that the time-limited offers in our design were relatively attractive mirrors the logic behind the use of such offers in retail markets. If, after rejecting a time-limited offer, a consumer can continue to search in a market in which free-recall offers exist, she is unlikely to be induced to accept a time-limited offer unless she perceives it to be attractive relative to the distribution of offers in the market as a whole. We used equal numbers of free-recall and time-limited tasks because we wanted the appearance of time-limited offers to be relatively infrequent and unpredictable. The presence of red cards in the free-recall tasks ensured that free-recall and time-limited tasks had exactly the same ex ante distribution of offers, made salient to subjects that favourable distributions of offers (i.e. offers on red cards) were not necessarily time-limited, and enabled us to identify the effect of identifying a card as red relative that specifically of it being time-limited.

4.3. Lottery tasks

As explained in Section 4.1, in Part 2 each subject faced fifteen lottery tasks, each requiring a choice between two lotteries. No time constraints were imposed on the choice itself. The parameters of these tasks were set separately for each subject, so that for each subject there was a one-to-one correspondence between the fifteen lottery tasks and the fifteen time-limited tasks from Part 1.¹² In defining this correspondence, we consider two alternative ways in which the subject might have responded to the relevant price search task. (These are not the only possible responses, but all other responses would be non-rational according to the criteria presented in Section 4.2.) *Response 1* is to accept the time-limited offer. *Response 2* is to reject the time-limited offer, not to accept any offer until all have been revealed, and then to choose the lowest free-recall offer. Lotteries 1 and 2 represent the distributions of earnings implied by Responses 1 and 2 respectively. That is, the certainty offered by Lottery 1 represents the earnings that the subject would have made in the matched price search task, had she accepted the time-limited offer. The risky prospect offered by Lottery 2 represents the distribution of possible earnings for the subject, conditional on the information that would have been available to her when the time-limited offer was visible, had she rejected that offer, waited until all free-recall offers had been revealed, and then chosen the lowest of these. Notice that the specification of each lottery task varied only according to the randomly-determined parameters of the corresponding time-limited task; it was independent of the subject's behaviour in that task. Notice also that if a price search

¹² Each subject faced the lottery tasks in the same order as she had faced the corresponding time-limited price search tasks. Recall, however, that the latter tasks were randomly distributed among the thirty price search tasks.

Lottery Task 3

Lottery 1

8.25

(100%)

Lottery 2

7.73	7.73	7.98	8.87	9.65
------	------	------	------	------

(20%) (20%) (20%) (20%) (20%)

Please indicate which lottery you want to choose:

choose Lottery 1
 choose Lottery 2

Next task

Fig. 2. Screen shot of a typical lottery task.

task is non-consequential, the corresponding lottery task is one in which Lottery 1 is stochastically dominated by Lottery 2. Conversely, if there is stochastic dominance in the lottery task, the price search task is non-consequential.¹³

For example, consider the price search task that generates the screenshot shown in Fig. 1. In this task, the time-limited price is 1.75EP. Since each task had a budget of 10EP, Response 1 gives earnings of 8.25EP with certainty. Thus, the outcome of the corresponding Lottery 1 is 8.25EP. The earnings from Response 2 are given by 10EP minus the lowest of the five free-recall prices. Two of these prices (2.27EP and 4.09EP) are already known; each of the others will be a random draw from the set {0.00EP, 0.01EP, ..., 10.00EP}. Thus, earnings from Response 2 are described by a well-defined probability distribution over {0.00EP, 0.01EP, ..., 10.00EP}. The corresponding Lottery 2 is a discrete approximation of that distribution. For each quintile of the actual distribution of Response 2 earnings, this approximation preserves the expected value of earnings conditional on earnings falling into that quintile. It therefore also preserves the unconditional expected value of earnings. Fig. 2 shows a screenshot of the lottery task that corresponds to the price search task of Fig. 1. It should be apparent from a comparison of the two figures that, although the two tasks are conceptually connected, there is no superficial resemblance between them that might facilitate memory connection or between-tasks learning.¹⁴

Each subject completed all fifteen lottery tasks before receiving any information about the actual outcome of any Lottery 2. At the end of the experiment, one lottery task was selected at random for payment. If the subject had chosen Lottery 2 in this task, the computer selected one of the five outcomes of that lottery at random, and this outcome determined the subject's earnings from Part 2. As far as Part 2 is concerned, the only difference between the two treatments occurred at this stage, and then only if the subject had chosen Lottery 1. In the No Feedback treatment, Lottery 2 was not resolved in

¹³ This statement is subject to a technical qualification. Because each Lottery 2 is a discrete approximation to the true distribution of possible earnings in the relevant price search task rather than that distribution itself, there is not an exact correspondence between non-consequential problems in price search tasks and stochastic dominance in lottery tasks.

¹⁴ Zizzo (2005) contains a summary of empirical evidence demonstrating how transfer of knowledge is easy when superficial features match but hard when they do not.

this case. In the Regret Feedback treatment, the computer determined an outcome for that lottery and this was displayed on the subject's screen. In each treatment, the Part 2 instructions made clear what information would be provided at the end of the experiment. This difference between the two treatments maintained the correspondence between price search tasks and lottery tasks.

4.4. Implementation

The experiment was conducted in fourteen sessions at the Centre for Behavioural and Experimental Social Science laboratory at the University of East Anglia during the summer of 2014. The experiment was programmed and conducted with the experimental software z-Tree (Fischbacher, 2007). Participants were recruited from the general university population by email, using the Hroot online recruitment system (Bock et al., 2014). Altogether there were 209 participants (101 male and 108 female), of whom 105 took part in the No Feedback treatment and 104 in the Regret Feedback treatment. Most of the participants were students from a wide range of academic disciplines and with an age range from 18 to 55. The experiment lasted about 50 minutes. Payments to participants ranged from £6.31 to £11.82, with an average of £10.43.

5. Results

5.1. Subjects' understanding of tasks

We begin by checking that subjects' behaviour showed an acceptable level of understanding of the tasks. In this section, we give an overview of the relevant findings; full details are given in Appendix 2.

One simple test is whether subjects chose dominated options. We will say that a subject chose a dominated option in a price search task if, at the moment at which she accepted an offer, another offer with a strictly lower price was visible and available. In the 6270 (= 209 × 30) responses to price search tasks, there were only 107 cases (1.7 per cent) in which dominated offers were chosen. This percentage was exactly the same for 4 s and 12 s tasks. Of the 3135 (= 209 × 15) cases of lottery tasks, 595 were choices between stochastically dominating and stochastically dominated lotteries; the dominated lottery was chosen in only 22 of these cases (3.7 per cent). These data suggest that subjects had a basic understanding of both types of task.

In the price search tasks, as explained in Section 4.2, a subject incurred no money cost and, at most, only minimal time cost by delaying a choice of a free-recall offer. We define a response to a price search task as *impatient* if the subject accepted a free-recall offer (whether on a red or blue card) before all offers had been revealed, and if that offer had a non-zero price and was not dominated. One might reasonably claim that impatient responses are not fully rational. In fact, 9.5 per cent of all responses in time-limited tasks (7.9 per cent at 4 s, 11.0 per cent at 12 s) and 23.3 per cent of all responses in free-recall tasks (19.8 per cent at 4 s, 26.7 per cent at 12 s) were impatient. Since choices of time-limited offers cannot be impatient by definition, the difference between the two types of tasks is unsurprising, but the high rate of impatience in free-recall tasks might suggest misunderstanding. However, a more detailed analysis reveals that, even when making impatient responses, most subjects terminated the search process only when they were able to accept very favourable offers. For any given subject who made an impatient response at any given stage of a free-recall task, we define the *loss due to impatience* as the difference between (i) the price the subject actually paid and (ii) the expected value of the price she would have paid, had she waited until all six offers had been revealed and then chosen the lowest price.¹⁵ The average loss per instance of impatience was only 0.13EP when the time interval was 4 s and 0.18EP when it was 12 s. For comparison, a subject who chose between the six offers at random would incur an expected loss of 3.89 EP.

We suggest that observed impatience is best explained as the behaviour of subjects who understood the tasks but used satisficing heuristics that accepted free-recall offers that were 'good enough' (Simon, 1957; Gigerenzer and Goldstein, 1996). There are various possible explanations of why, having rejected the time-limited offer, these subjects used such heuristics rather than inspecting all free-recall offers and choosing the best. They may not have realised that cutting short a search in one task would be unlikely to reduce the total length of the experiment for them. Or they may have felt that time spent waiting for other subjects to finish was less burdensome than time spent comparing offers. Or they may have felt social pressure not to delay the experiment for other subjects. Or they may simply have been impatient in the ordinary sense of the word – wanting to move on to the next task, without thinking about later consequences. But the low value of losses due to impatience suggests that, however it was caused, impatience had only a marginal effect on the outcomes of search, and that subjects placed low monetary valuations on the opportunity cost of time spent inspecting offers.

5.2. The data to be analysed

From now on, we will be concerned only with subjects' decisions between accepting and rejecting time-limited offers in consequential tasks, and with subjects' choices in the corresponding lottery tasks. Thus, the price search choice data that we analyse exclude the 19 per cent of cases in which the time-limited price was greater than or equal to the lowest preceding

¹⁵ This expected value is conditional on the subject's information at the relevant stage. The derivation of such expectations is explained in Appendix 3.

free-recall price. These data also exclude the very few cases in which a subject accepted a free-recall offer before seeing the time-limited offer.¹⁶ To maintain a one-to-one relationship between observations from price search tasks and lottery tasks, we also exclude the data from the corresponding lottery tasks. (By virtue of the principles used to construct lottery tasks, almost all of the excluded lottery tasks were ones in which Lottery 1 was stochastically dominated by Lottery 2.)

In testing Hypotheses 1 to 4, we compare the frequencies with which time-limited offers were chosen in different treatments or with different time constraints. Recall that, in each time-limited task faced by each subject, the price on each of the six cards was determined randomly, as was the position (first, second or third) of the time-limited offer in the sequence of cards. Therefore, the consequential tasks faced by a subject often included tasks in which the time-limited offer was worth much more than the expected value of the best free-recall price, and/or tasks in which it was worth much less. In comparing responses across treatments, it is important to control for this variation.

We do this by calculating, separately for each subject and for each relevant task, the expected value of the lowest free-recall price, conditional on the subject's information at the stage at which the time-limited offer appeared – that is, conditional on the actual values of offers that had already been revealed but using *ex ante* expectations of the others. From this expected value we subtract the actual time-limited price, to arrive at the *expected value difference* for that subject/task combination. (The derivation of expected value difference is explained in more detail in Appendix 3). The zero point on the scale of expected value difference corresponds with cases in which a fully rational, risk-neutral expected-utility-maximising individual would be indifferent between accepting and rejecting a time-limited offer. If the expected value difference is positive (negative), such an individual would accept (reject) the time-limited offer. In our analysis, however, this variable serves primarily as a one-dimensional summary statistic of the net advantage of accepting rather than rejecting a time-limited offer. A subject's response to any given consequential time-limited task can be summarised by the expected value difference in that task and her decision to accept or reject.

The calculation of expected value difference implicitly assumes that subjects do not accept (non-zero) free-recall offers until they have seen all six offers. If, to the contrary, some subjects are impatient and are able to predict their own impatience, expected value difference somewhat understates the net advantage of accepting a time-limited offer. The extent of this understatement corresponds with the expected value of 'losses due to impatience', as defined in Section 5.1. In fact, as shown in that section, these losses are very small. Thus, it is reasonable to use expected value difference as a measure of net advantage.

For any given category of time-limited tasks and aggregating across all subjects, we can investigate the relationship between the expected value difference and the relative frequency with which the time-limited offer is accepted. It is natural to expect that different subjects will have different attitudes to risk, and that the attitudes to risk revealed by any given subject will be subject to stochastic variation. Thus, one would expect this relationship – the *choice frequency function* – to be upward-sloping, with the S-shape that is characteristic of logistic functions. Throughout the paper, we show Lowess-smoothed plots of choice frequency functions. To provide a convenient summary statistic of the attitudes to risk revealed in any given category of time-limited task, we estimate the choice frequency function for that category using a logit model with subject-level clustering and then report the estimated probability with which the time-limited offer is chosen when the expected value difference is zero. We will call this the *standardised probability* of the choice of the time-limited offer in the relevant category of tasks.

5.3. Feedback effects

Fig. 3 shows Lowess-smoothed choice frequency functions for the No Feedback and Regret Feedback treatments. It is clear from this diagram that, contrary to Hypothesis 1, time-limited offers were slightly *less* likely to be chosen in the No Feedback treatment. The standardised probability of choosing the time-limited offer is 0.590 for the No Feedback treatment and 0.645 for the Regret Feedback treatment. To test whether behaviour in the two treatments is significantly different, we estimate a logit model with a dummy variable to represent the differential effect of the Regret Feedback treatment. The coefficient of this variable is positive (contrary to Hypothesis 1) but not significant ($p = 0.104$) (see Appendix 4, Section 1). Given our sample size, this test can pick up an effect size (i.e. between-treatment difference in standardised probabilities) of approximately 0.07.¹⁷ There was no significant interaction between regret-related feedback and time constraint in determining subjects' decisions (see Appendix 4, Section 3). As far as the hypothesis that we set out to test is concerned, our first main result is clear:

Result 1: There is no support for Hypothesis 1, i.e. the hypothesis that time-limited offers are more likely to be chosen when feedback is absent than when it is present.

In Sections 5.4 and 5.5, we pool the data from the two treatments. Because the treatments are identical in all respects other than the provision of feedback, this pooling cannot induce systematic biases in our results; it merely introduces a small amount of additional noise.

¹⁶ 14 of the 3135 responses to time-limited tasks (0.04 per cent) were of this kind.

¹⁷ This statistic is calculated for a two-tail test at 5 per cent significance and a power of 0.8 using Optimal Design Software Version 3.0 (Spybrook et al., 2011). See Appendix 4, Section 2.

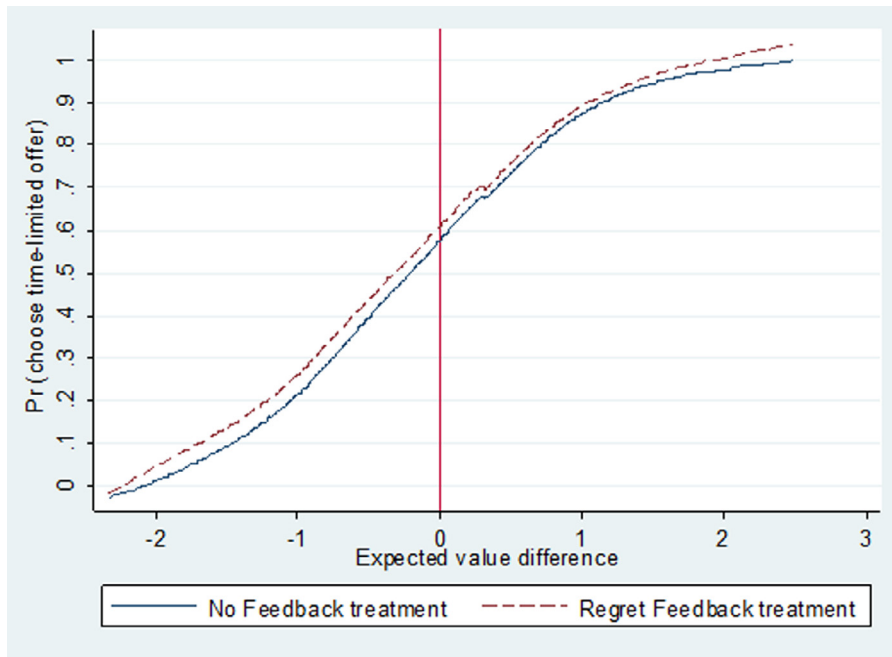


Fig. 3. Choice frequency functions for no feedback and regret feedback treatments.

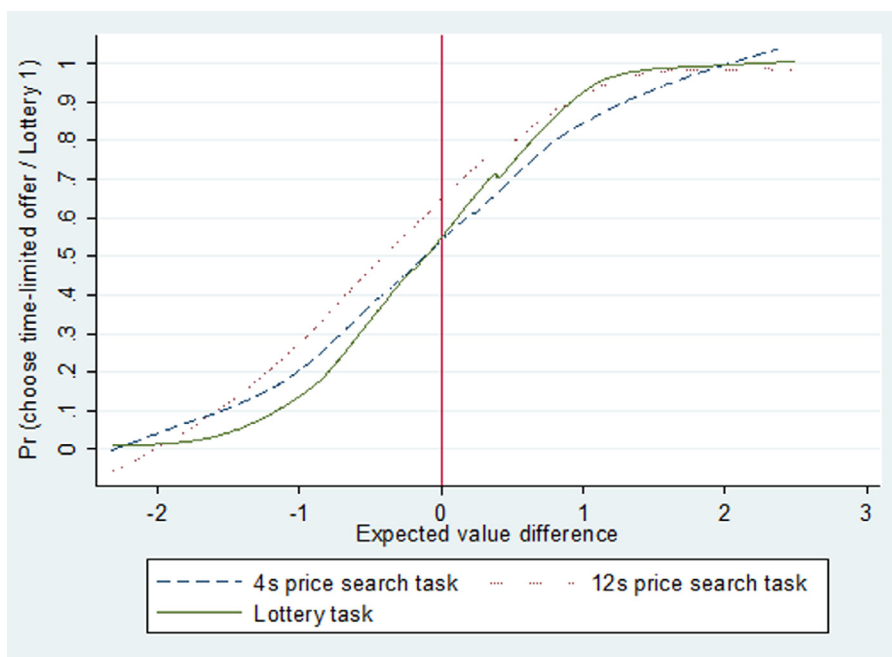


Fig. 4. Choice frequency functions for 4 s and 12 s price search tasks and lottery tasks.

5.4. Time constraints

We now use the same method as in Section 5.3 to test Hypothesis 2 about the effect of time constraints. Fig. 4 plots Lowess-smoothed choice frequency functions for the 4 s and 12 s time-limited tasks. (The choice frequency function for lottery tasks will be discussed later.)

It is immediately obvious that, in direct opposition to Hypothesis 2, time-limited offers are *less* likely to be chosen in 4 s tasks than in 12 s tasks. The standardised probability of choosing the time-limited offer is 0.551 for 4 s tasks and 0.691 for 12 s tasks. To test whether behaviour is significantly different at the two levels of time constraint, we estimate a logit

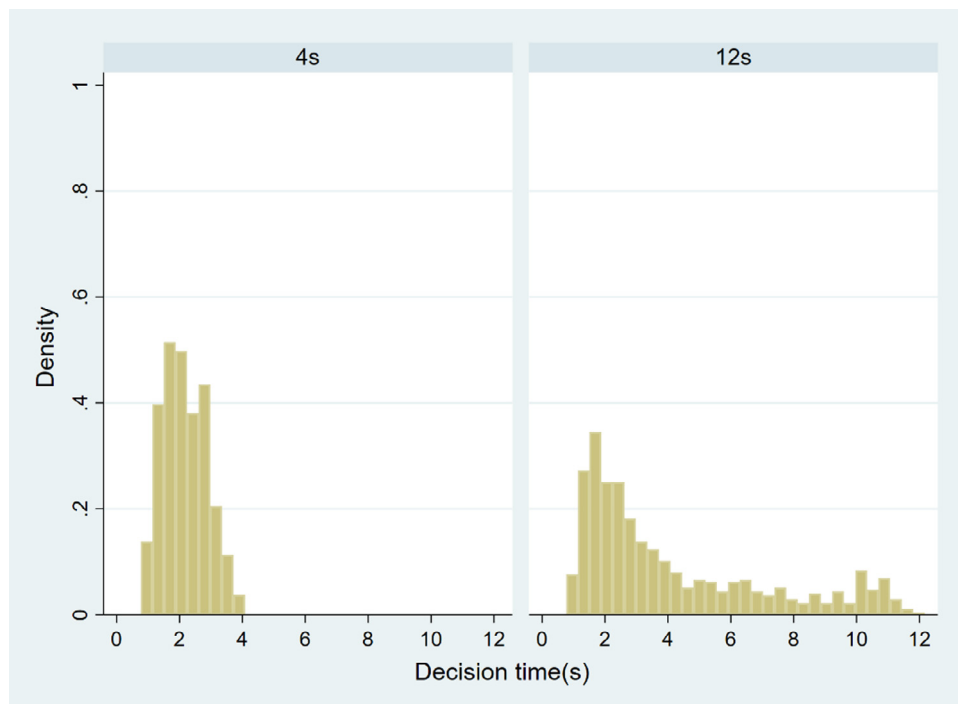


Fig. 5. Distributions of decision times for 4 s and 12 s price search tasks.

model with subject-level clustering. A dummy variable takes the value 0 for 4 s tasks and 1 for 12 s tasks. The coefficient of this variable is positive and strongly significant ($p < 0.001$) (see Appendix 4, Section 4).

Choice frequency functions can also be used to show the degree of dispersion in responses, controlling for variation in expected value difference. It is clear from Fig. 4 that the degree of dispersion is similar at the two levels of time constraint. Thus, there is no evidence to suggest that individual decision-making was subject to more stochastic variation when the time available was shorter, as would be implied by many psychological theories of decision processes (e.g. Busemeyer and Townsend, 1993; Stewart et al., 2006). It is also useful to look at *decision times* for the acceptance of time-limited offers (i.e. time intervals between when a time-limited offer appeared and when the 'choose this offer' button was clicked). Recall that a rejection was recorded if the button had not been clicked at the deadline. If the deadline cut in while a subject was still thinking about her decision, unintended rejections might be recorded. If such outcomes were common, the distribution of decision times would be truncated at the deadline. Fig. 5 shows the distributions of decision times for 4 s and 12 s tasks. The bell-shaped 4 s distribution suggests that, even with this degree of time constraint, unintended rejections were rare. Notice also that a large proportion (62.9 per cent) of 12 s tasks was completed within 4 s. This suggests that 4 s was sufficient for many subjects to complete all the reasoning they wanted to do.

Another possibility is that the difference between the two choice frequency functions in Fig. 4 is an artefact of impatience. Recall that subjects were more impatient in 12 s tasks than in 4 s tasks, and that, if subjects are impatient, our measure of expected value difference understates the net advantage of choosing the time-limited offer and does not take account of any perceived opportunity cost of time spent inspecting offers. Thus, some part of the difference between the two functions may be due to greater understatement of net advantage in the 12 s case. But, as explained in Sections 5.1 and 5.2, the low observed value of losses due to impatience suggests that any such understatement is relatively small, and that the perceived opportunity cost of time is low, irrespective of the causes of impatience. In Section 6, we report a second experiment that allows us to test Hypotheses 2 and 3 while implementing stronger controls for the effects of impatience. The data from that investigation confirm the second main result from the main experiment:

Result 2: Contrary to Hypothesis 2, time-limited offers are less likely to be chosen when the time constraint is tighter.

For both 4 s and 12 s tasks, the standardised probability of choosing the time-limited offer was significantly greater than 0.5 ($p = 0.028$ for 4 s tasks, $p < 0.001$ for 12 s tasks). Hence:

Result 3: Consistently with Hypothesis 3, subjects' choices between accepting and rejecting time-limited offers reveal an overall tendency towards risk aversion. This tendency towards risk aversion is stronger when the time constraint is looser.

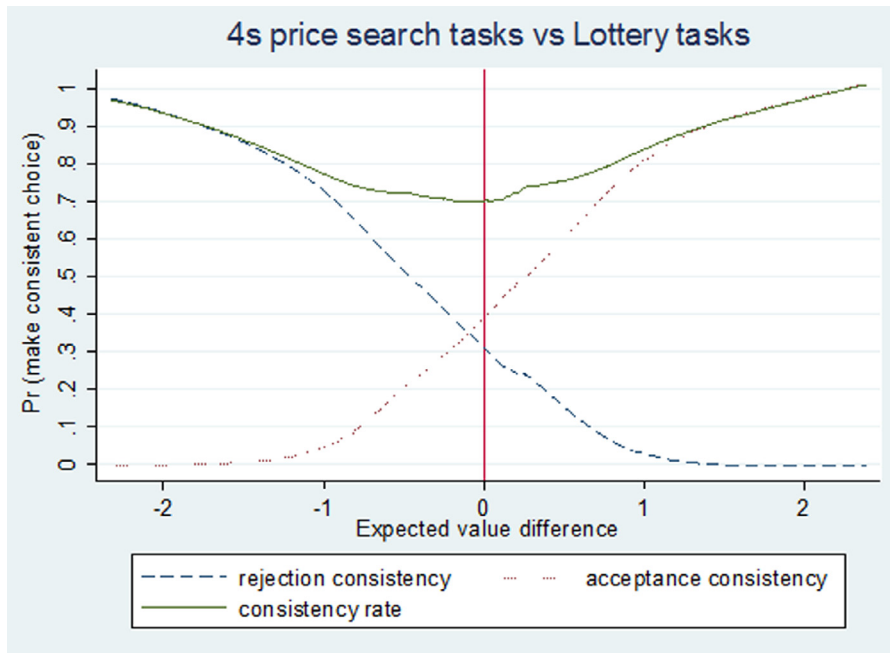


Fig. 6. Consistency of individual choices between 4 s price search tasks and lottery tasks.

5.5. Comparisons between price search tasks and lottery tasks

We now consider subjects' responses to the lottery tasks. Analogously with our treatment of price search tasks, we define expected value difference for any given lottery task as the (certain) value of Lottery 1 *minus* the expected value of Lottery 2. Recall that, for each <subject, task> pair that is included in our analysis, there is a corresponding lottery task. Payoffs in this lottery task correspond with potential earnings in the price search task. Our definition ensures that every lottery task has the same expected value difference as the price search task to which it corresponds. Aggregating over all subjects and all (non-excluded) lottery tasks, we can plot the relative frequency of Lottery 1 choices as a function of the expected value difference.¹⁸ Fig. 4 (see Section 5.4 above) shows the Lowess-smoothed choice frequency function for lottery tasks, together with the corresponding functions for the 4 s and 12 s price search tasks.

The standardised probability of choosing Lottery 1 is 0.577. This is significantly different from 0.5 ($p < 0.001$), indicating an overall tendency towards risk aversion. To test whether attitudes to risk revealed in the lottery tasks are significantly different from those revealed in the price search tasks, we estimate a logit model with subject-level clustering. We use two dummy variables. One takes the value 1 in 4 s price search tasks, 0 otherwise; the other takes the value 1 in 12 s price search tasks, 0 otherwise. The 4 s dummy variable is negative but not significant ($p = 0.959$); the 12 s dummy variable is positive and highly significant ($p < 0.001$) (see Appendix 4, Section 6). Since lottery tasks were faced later in the experiment than price search tasks, differences in risk aversion between the two types of tasks might in principle be an artefact of gradual changes in risk aversion over the course of the experiment. But we found no trend in risk aversion over the thirty price search tasks (see Section 5.6). We conclude:

Result 4: When the time constraint is loose, decisions about accepting or rejecting time-limited offers are more risk-averse than are binary choices between certainties and lotteries. When the time constraint is tight, there is no significant difference between attitudes to risk in the two types of decision problem.

Notice that the tests that support Result 4 do not take account of differences in dispersion of attitudes to risk. It is clear from inspection of Fig. 4 that there is less such dispersion in lottery tasks than in price search tasks. This is perhaps not surprising. In a lottery task, the probability distributions of payoffs for the two options are described explicitly to subjects. In a price search task, in contrast, it is a difficult problem to work out from first principles the distribution of payoffs implied by a rejection of the time-limited offer. If subjects' beliefs about this distribution are wholly or partly derived from experience, or if they use experience-based decision heuristics which bypass the formation of such beliefs, the random elements and time lags of the learning process will introduce additional noise into their responses.

¹⁸ If the lottery choice data are disaggregated between the No Feedback and Regret Feedback treatments, the choice frequency functions for the two treatments are almost identical. Plots of these functions are shown in Appendix 4, Section 5.

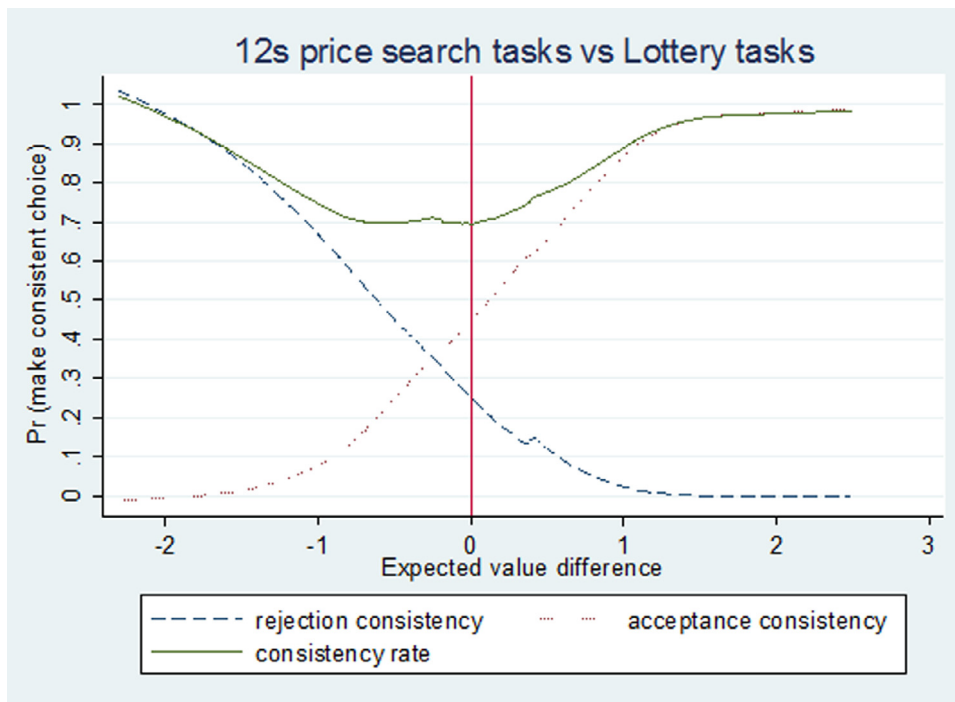


Fig. 7. Consistency of individual choices between 12 s price search tasks and lottery tasks.

Fig. 4 shows a broad pattern of similarity between aggregate responses to time-limited tasks (particularly at 4 s) and aggregate responses to lottery tasks. But it is also useful to investigate consistency between responses to the two types of task at the level of the individual subject. For each subject and for each consequential time-limited price search task faced by that subject, we can compare her response to that task with her response to the matched lottery task. We will say that these responses show *acceptance consistency* if the subject chose the time-limited offer in the price search task and Lottery 1 in the lottery task, and *rejection consistency* if the time-limited offer is rejected and Lottery 2 is chosen. Aggregating across subjects and task pairs, the relative frequency of the two types of consistency can be plotted as functions of expected value difference.

Fig. 6 shows the Lowess-smoothed frequency functions for the two types of consistency for 4 s tasks. The sum of these relative frequencies (the *consistency rate*) is also plotted. Fig. 7 shows the corresponding plots for 12 s tasks.

In interpreting consistency rates, it is useful to have a benchmark. Previous experimental research has generated data about how subjects respond when *exactly the same* problem of choosing between two lotteries is faced twice in the same experiment. These choice problems have normally been designed so that neither option is obviously better than the other. Consistency rates in such problems have typically been found to lie in a range from 70 to 82 per cent.¹⁹ Intuitively, one might expect much less consistency when, as in our experiments, the two decision problems are framed very differently. A further reason for expecting inconsistencies arises from the distinction between *decisions from description* (that is, problems of choice under uncertainty in which decision-makers are given explicit information about probabilities) and *decisions from experience* (problems that are repeated several times and for which probabilities have to be inferred from experience). Our lottery tasks are clearly decisions from description, while (for reasons explained in Section 3) subjects might treat our price search tasks as decisions from experience. It is known that behaviour is systematically different in the two environments. In particular, when individuals are making decisions by description, they tend to over-weight small probabilities (Kahneman and Tversky, 1979), but when decisions are made by experience, small probabilities tend to be under-weighted (Hertwig et al., 2004).

Viewed against the benchmark of repeated identical decision problems, consistency rates in our experiment are remarkably high. In both 4 s and 12 s tasks, the consistency rate reaches a minimum in the region of expected value difference at which each option is chosen with approximately the same frequency; in both cases, the minimum value of this rate is just below 70 per cent. Hence:

¹⁹ Loomes et al. (2002) report an experiment with a consistency rate of 82 per cent and cite four other studies which found consistency rates between 70 and 80 percent.

Table 1
Regression results using expected value difference.

	Overall (1)		NF (2)		RF (3)	
	β	ME	β	ME	β	ME
Treatment	0.329 (0.200)	0.078 (0.047)				
Expected value difference	2.399*** (0.104)	0.572*** (0.026)	2.692*** (0.168)	0.662*** (0.043)	2.139*** (0.128)	0.496*** (0.030)
Available time	0.692*** (0.130)	0.164*** (0.030)	0.634** (0.195)	0.154** (0.047)	0.715*** (0.175)	0.164*** (0.039)
Good offers rejected	2.105*** (0.272)	0.502*** (0.065)	2.960*** (0.450)	0.728*** (0.111)		
Good offers seen					1.430*** (0.351)	0.331*** (0.082)
Period	0.003 (0.008)	0.001 (0.002)	0.005 (0.012)	0.001 (0.003)	0.004 (0.011)	0.001 (0.002)
Constant	-0.548** (0.201)		-0.728** (0.276)		-0.386 (0.245)	
Observations	2456		1246		1210	
LR Chi-Square	536.210		256.673		280.878	
Prob > Chi-Square	0.000		0.000		0.000	
Baseline predicted probability	-0.794		-1.061		-0.643	

Notes: * 10% level, ** 5% level, *** 1% level. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. We used panel data to estimate all these models. The data used to estimate model 2 contains 1246 observations from 105 subjects in the No Feedback treatment. The data used to estimate model 3 contains 1210 observations from 104 subjects in the Regret Feedback treatment. For each model, the left column contains coefficients, and the right column report marginal effects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

Result 5: At the level of individual subjects, the rate of consistency between time-limited offer choices and corresponding lottery choices is similar in magnitude to previously-observed consistency rates between identical lottery choice problems.

This degree of consistency between price search tasks and lottery tasks suggests that, for a typical subject, both tasks tapped into some common body of attitudes to risk: preferences were not simply constructed in response to specific decision problems. To put this another way, individual-level differences in willingness to accept time-limited offers generally reflected differences in attitudes to risk, as revealed in choices over binary lotteries. This is consistent with Brown et al.'s (2018) finding that risk aversion (as measured by questionnaire responses) was positively correlated with the acceptance of time-limited offers that would have been rejected by risk-neutral rational agents (see Section 2). Further, and in line with the findings reported in Section 5.4, the observed consistency between search and lottery tasks suggests that subjects' capacity to make considered decisions about such offers was not greatly impaired by the tight time constraint of the 4 s tasks. Nevertheless, they made systematically different decisions when time constraints were relaxed.

5.6. Learning

Finally, we consider how subjects' responses evolved over the course of the experiment. To investigate this, we regress subjects' decisions in time-limited tasks on a set of explanatory variables, some of which capture possible learning mechanisms. All the regressions we report are estimations of logit models in which subject-level random effects are controlled.

Results for our three main models are presented in Table 1. In these models, the dependent variable is the probability that a subject chooses the time-limited offer in a consequential time-limited task. Because subjects in the No Feedback and Regret Feedback treatments received different information, the distinction between these treatments is potentially significant in explaining learning. Model 1 uses data from all time-limited tasks; Models 2 and 3 respectively use only data from the No Feedback and Regret Feedback treatments.

Treatment in Model 1 is a dummy that takes the value 1 if subjects are in the Regret Feedback treatment. *Expected value difference* is as defined in Section 5.2. *Available time* is a dummy which takes the value 1 in 12 s tasks. *Period* represents the position of the task in the sequence of price search tasks faced by the subject; 1 is the first task and 30 is the last. This variable picks up overall trends in learning.

We define two additional variables to pick up the effects of experience of the realisations of offer values. For each consequential time-limited task faced by each subject, we define the time-limited offer as *good* if its price is strictly less than

the lowest price in the set of five free-recall offers and *bad* otherwise.²⁰ Notice that the good/bad distinction is based on ex post realisations of offer values, not ex ante expectations, and is unaffected by whether or not the time-limited offer was chosen. At the end of each time-limited task in the Regret Feedback treatment, subjects always knew whether the time-limited offer was good or bad. For subjects in this treatment, experiential learning might be mediated by memories of good and bad offers observed in previous tasks. Such memories might be interpreted as experiences of regret (if a good offer was rejected or a bad offer accepted) or rejoicing (in the opposite cases). Alternatively, they might be interpreted as encoding information about the distribution of the difference in value between the time-limited offer and the best free-recall offer. The variable *Good offers seen* is defined (for Model 3 only) as the proportion of previous time-limited tasks in which the time-limited offer was good. (In a subject's first time-limited task, this variable takes the value zero.)

For Models 1 and 2, we are forced to use a less clean measure of a subject's experience of good and bad offers. To maintain as close a parallel as possible with *Good offers seen*, we define *Good offers rejected* as the number of previous tasks in which good time-limited offers were not chosen as a proportion of the total number of previous tasks in which time-limited offers were not chosen. (In the first task in which a subject did not choose the time-limited offer, this variable takes the value zero.) Although subjects who made impatient decisions in the No Feedback treatment did not see all five free-recall offers, the value of the best of these offers is a good approximation to the value of the best free-recall offer actually seen by the subject (see Section 5.1). A further confound needs to be taken into account. For subjects who (independently of experience) have a relatively strong propensity to accept time-limited offers, the proportion of rejected offers that are found to be good will be relatively low. This mechanism induces a tendency for the choice of time-limited offers to vary negatively with *Good offers rejected*. Thus, finding a positive effect would be particularly strong evidence of reinforcement or regret-based learning.

Table 1 shows the results of these regression models, including the coefficients and the marginal effects. In all three regressions, unsurprisingly, the coefficients for *Expected value difference* are positive and highly significant ($p < 0.001$ in all cases). *Available time* has a positive and highly significant effect in all three regressions, in line with the findings reported in Section 5.4. In Model 1, the coefficient for *Treatment* is positive but not quite significant at the 10 per cent level ($p = 0.100$), paralleling the findings reported in Section 5.3. *Period* is insignificant in all three regressions, indicating the absence of any trend in the probability with which time-limited offers are chosen.²¹

In Model 3, *Good offers seen* has a positive and highly significant effect, indicating reinforcement or regret-based learning. Similarly, *Good offers rejected* has a positive and highly significant effect in Models 1 and 2. The parallel between the effects of these two experience variables strongly suggests that, despite the confounds discussed in the previous paragraph, reinforcement or regret-based learning is at work in both treatments. To investigate whether this learning mechanism was affected by the tightness of time constraints, we tried adding variables to pick up interactions between *Available time* and *Good offers seen* (in Model 3) or *Good offers rejected* (in Models 1 and 2). We found no significant interaction effects (see Appendix 4, Section 8).

In our analysis so far, we have used *Expected value difference* as a summary measure of the net advantage to be expected from choosing rather than rejecting the time-limited offer. However, it is also useful to investigate how subjects used specific items of information about the offers in a task. We therefore estimated three additional models in which *Expected value difference* was replaced by three *offer variables* that were directly observed by subjects: the price of the time-limited offer (*Red card*), the position of that offer in the sequence of cards (*Position*), and the lowest free-recall price revealed before the appearance of the time-limited offer (*Best blue card*). Because *Best blue card* is undefined if the time-limited offer is the first in the sequence, we restrict the analysis to cases in which that offer appeared second ($Position = 0$) or third ($Position = 1$). Notice that the distribution of outcomes resulting from the rejection of a time-limited offer is unambiguously better, the lower the values of *Best blue card* and *Position* (i.e. the more free-recall offers remain to be revealed). These models are reported in Table 2.

In all three models, *Red card*, *Position* and *Best blue card* have their expected signs. The first two variables are highly significant in all three models. *Best blue card* is significant at the 5 per cent level in Model 1 (using data from both treatments) and Model 3 (using only data from the Regret Feedback treatment). As in the models reported in Table 1, *Good offers seen* and *Good offers rejected* are highly significant; there are no significant interactions between these variables and *Available time* (see Appendix 4, Section 9). Nor are there any significant interactions between the offer variables and *Available time* (see Appendix 4, Section 10).

6. Further evidence about the effects of time constraints

Readers of earlier versions of this paper have expressed surprise that we found time-limited offers to be less likely to be chosen when time constraints were tighter. To check that result and to gather more evidence about the effects of variation in time constraints, we ran a second experiment. Except for two points of difference, that experiment (*Experiment 2*) had

²⁰ The random realisation of the value of each card in each task was determined before the subject made any decisions about turning over cards. Thus, 'good' and 'bad' cards are well-defined even if the subject never saw them turned over.

²¹ As an additional check, we compared the Lowess-smoothed choice frequency functions for the choice of the time-limited offer (i) using data only from the first fifteen price search tasks faced by each subject and (ii) using only the last fifteen price search tasks. We made these comparisons separately for 4s and 12s tasks. In each case, there was no significant difference between the functions. See Appendix 4, Section 7.

Table 2
Regression results using offer variables.

	Overall		NF		RF	
	(1)		(2)		(3)	
	β	ME	β	ME	β	ME
Treatment	0.307 (0.224)	0.064 (0.047)				
Red card	-2.146*** (0.128)	-0.452*** (0.028)	-2.461*** (0.215)	-0.514*** (0.050)	-1.895*** (0.155)	-0.403*** (0.033)
Position	0.829*** (0.178)	0.175*** (0.038)	0.762*** (0.264)	0.159*** (0.056)	0.860*** (0.237)	0.183*** (0.050)
Best blue card	0.097** (0.040)	0.020** (0.008)	0.074 (0.061)	0.015 (0.013)	0.129** (0.052)	0.028** (0.011)
Available time	0.673*** (0.169)	0.143*** (0.036)	0.681*** (0.255)	0.143*** (0.054)	0.647*** (0.226)	0.139*** (0.049)
Good offers rejected	1.866*** (0.338)	0.393*** (0.071)	2.780*** (0.566)	0.580*** (0.117)		
Good offers seen					1.232*** (0.461)	0.262*** (0.098)
Period	-0.001 (0.010)	-0.000 (0.002)	-0.013 (0.016)	-0.003 (0.003)	0.011 (0.014)	0.002 (0.003)
Constant	1.190** (0.595)		2.089** (0.900)		0.469 (0.770)	
Observations	1444		738		706	
LR Chi-Square	292.027		134.373		155.455	
Prob > Chi-Square	0.000		0.000		0.000	
Baseline predicted probability	-0.197		-0.224		-0.212	

Notes: * 10% level, ** 5% level, *** 1% level. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. We used panel data to estimate all these models. The data used to estimate model 2 contains 1444 observations from 105 subjects in the No Feedback treatment. The data used to estimate model 3 contains 706 observations from 104 subjects in the Regret Feedback treatment. For each model, the left column contains coefficients, and the right column reports marginal effects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

exactly the same design as the original one (*Experiment 1*). The first difference was that, after the red card had been turned over to reveal its offer (whether time-limited or free-recall), the time interval before a blue card was turned over could be any of 2 s, 4 s, 8 s, 12 s or 18 s. Participants were not told which interval would apply until the red card was turned over. Thus, we were able to investigate five different degrees of time constraint, including the two that were used in *Experiment 1*.

The second difference was that, irrespective of the time interval for the red card, the time interval for blue cards (i.e. the time between a blue card being turned over and the turning-over of the next card) was always 4 s. This design feature provides a control for the potential confound of impatience, discussed in [Section 5.4.22](#). Recall that our object was to compare the effects of different levels of time constraint on decisions between accepting and rejecting time-limited offers in consequential tasks. The potential confound in *Experiment 1* was that, conditional on the rejection of a time-limited offer, the time required to search amongst later free-recall offers was greater in tasks in which the time constraint for the time-limited offer was looser. Combined with the assumption that subjects were impatient to move on to the next task, this difference in the time required to find the best free-recall offer is a possible artefactual explanation of why time-limited offers were chosen more frequently when time constraints were looser. In the new design, the alternative to accepting a time-limited offer was always the prospect of facing a known number of randomly-generated free-recall offers, appearing at 4 s intervals. Since the properties of this prospect were independent of the time constraint for accepting or rejecting the time-limited offer, differences between subjects' responses to different time constraints should not be affected by attitudes to time spent looking at free-recall offers. We chose a short blue-card interval to minimise demands on participants' patience, given that a rational subject would never want to choose a free-recall offer before seeing all six offers.

Even in the new design, of course, impatient subjects might perceive time spent *considering time-limited offers* as subjectively costly. In this sense, longer time constraints impose greater time costs. This is not a confound in itself: in any natural or experimental setting, relaxing a time constraint for a decision *must* provide an opportunity to divert time from some other use to the act of decision-making. In our design, however, a subject can reject a time-limited offer only by waiting for the next offer to appear. Thus, it might be argued that the interval between an impatient subject's (unobserved) mental act of rejecting a time-limited offer and the arrival of the next offer is a cost which varies with the time constraint. But if this effect is a confound, it is far smaller than the corresponding effect in *Experiment 1*. In any consequential task, it affects only the time-limited offer, rather than that offer plus three, four or five free-recall offers; and it applies only to a fraction

²² We thank an editor for suggesting this improvement to the original design.

Table 3
Standardised probabilities of choice of time-limited offer.

Task	Standardised probability:	
	Experiment 1	Experiment 2
Price search:		
2s	—	0.316
4s	0.551	0.523
8s	—	0.553
12s	0.691	0.701
18s	—	0.611
Lottery	0.577	0.589

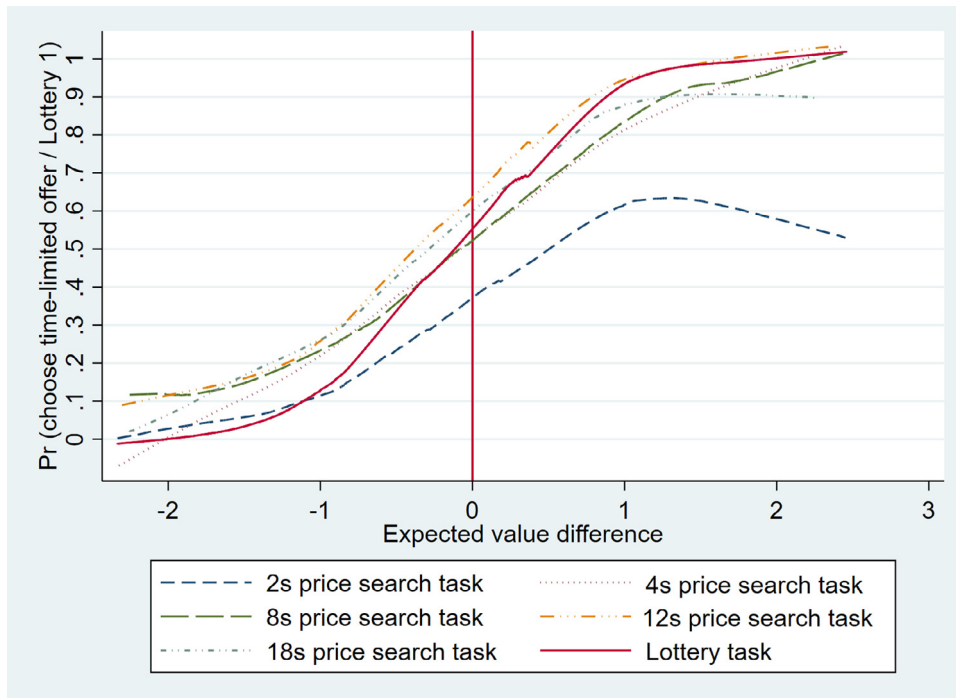


Fig. 8. Choice frequency functions for Experiment 2.

of the time that offer is displayed. If (as in fact we find) the elimination of 80 per cent or more of the putative confound were to have no material effect on our results, it would be reasonable to infer that those results were not an artefact of the remainder.

Experiment 2 was run in October 2018. The 136 participants were recruited from the same subject pool as was used for Experiment 1, except for the turnover in the university population between the two dates. No one participated in both experiments. In this section, we present only the most important results; more details are reported in Appendix 5.

The overall frequency of impatient responses was similar to that in the 4 s tasks of Experiment 1. As in Experiment 1, expected losses due to impatience were very small. Using the same tests as before (see Section 5.3), we again found no significant difference between the No Feedback and Regret Feedback treatments, and no interaction between feedback and time constraints (see Appendix 5, Section 1). The standardised probability of choosing the time-limited offer was almost the same in both cases (0.602 in the No Feedback treatment, 0.588 in the Regret Feedback treatment). This replication further increases our confidence in Result 1. As before, we pool the data from the two treatments when analysing issues other than feedback.²³

The main results of Experiment 2 in relation to time constraints and comparisons between price search and lottery tasks are summarised in Table 3 (which also shows the corresponding data from Experiment 1) and in the Lowess-smoothed choice frequency functions shown in Fig. 8.

²³ These probabilities are calculated after excluding 2 s tasks. Our reason for excluding these tasks is explained below.

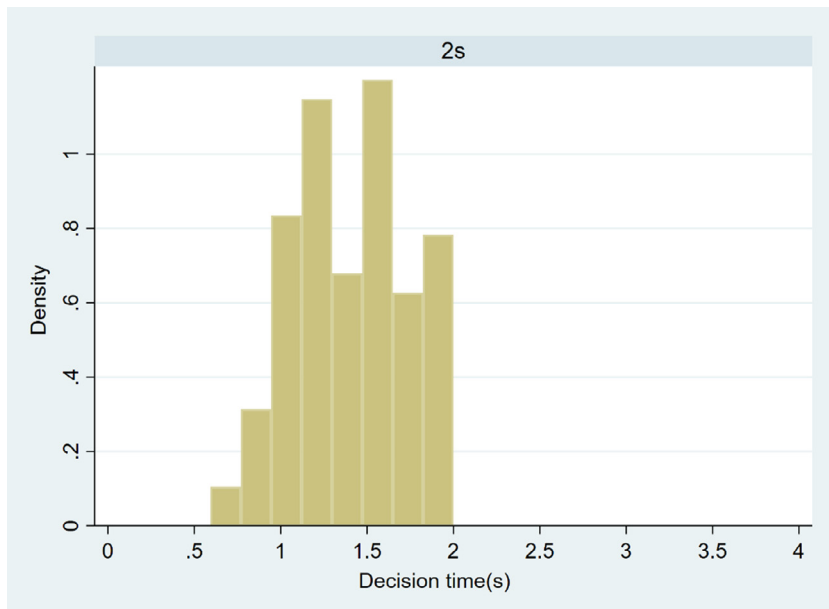


Fig. 9. Distribution of decision times for 2 s price search tasks in Experiment 2.

It is immediately obvious that there is a discontinuity between the 2 s tasks and the others. Table 3 shows that the standardised probability for the choice of the time-limited offer was only 0.316 in the 2 s task, compared with values in the range from 0.523 to 0.701 when the time interval was 4 s or greater. Fig. 8 shows that, in 2 s tasks, the frequency with which the time-limited offer was chosen remained below 0.65 even when the time-limited offer was very favourable. The distributions of decision times give further insight into this discontinuity. As in Experiment 1, the 4 s distribution is bell-shaped; at longer time intervals, the modal decision time is approximately 2 s, with a large proportion of decisions being made within 4 s (see Appendix 5, Section 2). In contrast, the 2 s distribution, shown in Fig. 9, is clearly truncated at the 2 s deadline. The most obvious interpretation is that, for a significant proportion of subjects, that deadline cut in while deliberation was still in progress. These subjects would be recorded as having rejected the time-limited offer, even though they had not consciously chosen to do so. We therefore exclude the 2 s tasks from analyses that treat the degree of time constraint as a continuous variable. Recall that the main purpose of Experiment 2 was to check our original finding that time-limited offers were chosen less frequently when time constraints were tighter. By excluding 2 s tasks, we screen out a confounding factor that would have favoured the original finding.

Because Experiment 2 had fewer participants than Experiment 1 and used five different time intervals, the number of observations for each interval considered separately is smaller than in the original experiment. One should therefore not read too much into data that is specific to any one interval; it is more useful to look for patterns in the data as a whole (after excluding 2 s tasks). The choice frequency functions for 4 s, 8 s, 12 s and 18 s show the same qualitative features as were found for 4 s and 12 s in Experiment 1. As in Experiment 1, standardised probabilities are greater than 0.5, indicating a tendency towards risk aversion. There is a clear tendency, visible both in the Table 3 data and the Fig. 9 graphs, for time-limited offers to be chosen more frequently when time constraints are looser. If the time interval is treated as a continuous variable, it has a positive and strongly significant effect on the probability of choosing the time-limited offer ($p=0.008$), as shown in the regression in Table 4. These findings corroborate Results 2 and 3.

As in Experiment 1, responses to lottery tasks were less dispersed than responses to price search tasks (as is immediately obvious from Fig. 9). There was partial corroboration for Result 4. As in Experiment 1, decisions were significantly more risk averse in 12 s price search tasks than in lottery tasks ($p=0.001$), but there were no significant differences at 4 s, 8 s or 18 s (see Appendix 5, Section 3). At the individual subject level, responses to price search tasks (at each of 4 s, 8 s, 12 s and 18 s) were remarkably consistent with responses to lottery tasks (see Appendix 5, Section 4). Using the new data, we ran the same regressions as reported in Tables 1 and 2 in Section 5.6, but with available time as a continuous variable rather than a dummy, and found essentially the same results (see Appendix 5, Section 5).

7. Discussion

In this section, we discuss three unexpected features of our results.

The first of these is the absence of a regret feedback effect. Previous experiments and surveys have found a tendency for individuals to choose options that reduce their exposure to regret (Zeelenberg et al., 1996; Zeelenberg and Pieters, 2004).

Table 4
Regression analysis for Experiment 2.

	(1)	
	β	ME
Expected value difference	1.642*** (0.102)	0.396*** (0.025)
Available time	0.032*** (0.012)	0.008*** (0.003)
Constant	0.047 (0.153)	
Observations	1337	
Pseudo R ²	0.350	
LR Chi-Square	258.679	
Prob > Chi-Square	0.000	
Baseline predicted probability	0.403	

Notes: * 10% level, ** 5% level, *** 1% level. Cluster-robust standard errors in parentheses.

The dependent variable in the model is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. *Available time* is a continuous variable. For the model, the left column contains coefficients, and the right column reports marginal effects. Results for the model are based on logit estimations with subject-level clustering.

Our No Feedback treatment was set up so that choosing a time-limited offer would eliminate the feedback that could otherwise induce regret, but we found no evidence that this contributed to the attractiveness of such offers.

One possible explanation of this difference in findings is that the regrets that our time-limited tasks were likely to induce were relatively mild.²⁴ In contrast, the experiments reported by Zeelenberg et al. involve hypothetical choices between binary lotteries in which the worse outcome is always winning nothing. The case studied by Zeelenberg and Pieters is particularly extreme: a person who chooses not to buy a ticket in the Dutch Postcode Lottery is exposed to the possibility of massive regret. Such extreme cases are obviously relevant in some natural-world cases. That said, if exposure to regret is measured in relative terms (i.e. based on the comparison in value between ‘what is’ and ‘what might have been’), the modest exposure to regret in our experiment is typical of a number of natural-world price search problems.²⁵ Although our subjects seem not to have been influenced by *anticipations* of regret, their willingness to accept time-limited offers was affected by their previous experience of offer realisations (as explained in Section 5.6). However, the absence of a regret feedback effect suggests that the mechanism by which experience was encoded and recalled was not driven by the aversive effects of regret.

The second unexpected feature of our results is the high degree of consistency, at both the aggregate and individual levels, between time-limited offer choices made even under tight time constraints and the corresponding lottery choices, for which there were no time constraints. This consistency is particularly noteworthy given the very different presentation of the two decision problems. It is also noteworthy as subjects’ decisions about time-limited offers were strongly influenced by realisations of offer prices in previous tasks. That influence suggests that price search tasks were treated as decisions from experience. In contrast, since no lottery risk was resolved until the end of the experiment, lottery tasks were necessarily decisions from description.²⁶

It is clear from the results reported in Section 5.6 that, in choosing whether or not to accept a time-limited offer, subjects took account of the value of that offer and of its position in the sequence of offers. The evidence also suggests that

²⁴ Consider any consequential task in which the time-limited offer was rejected and in which this offer was lower than the lowest free-recall offer. We define the associated ‘regret’ as the lowest free-recall offer minus the time-limited offer. Using Experiment 1 data only from cases that gave rise to regret, the mean (median) regret from rejecting time-limited offers was 1.13EP (0.70EP) in the Regret Feedback treatment and 1.46EP (1.07EP) in the No Feedback Treatment. In the Regret Feedback treatment, if the time-limited offer was accepted and if this offer was higher than the lowest free-recall offer, ‘regret’ is defined as the time-limited offer minus the lowest free-recall offer. In this case, the mean (median) regret was 0.93EP (0.75EP). In future research, it would be interesting to see whether questionnaire based measures of regret change as a result of the rejection of time-limited offers.

²⁵ There is some evidence that people tend to over-predict the intensity of future regrets (Gilbert et al., 2004). Thus, repetition of tasks (as in our experiment) might be expected to lead to a progressive attenuation of regret effects. However, we found no significant trends in our data (see Section 5.6). Engelbrecht-Wiggans (1989) finds evidence that, in a first-price sealed-bid auction, ‘winner’s regret’ (related to an error of commission) is more painful than ‘loser’s regret’ (related to an error of omission). Arguably, not accepting an time-limited offer is an omission. However, the best evidence of feedback-conditional regret comes from postcode lotteries, where the relevant regret is clearly about an omission. In future research, it would be interesting to see whether questionnaire based measures of regret change as a result of the rejection of time-limited offers.

²⁶ To check whether subjects’ lottery choices might have been influenced by realisations of offer values in the preceding price search tasks, we estimated logit models in which the dependent variable was the choice of Lottery 1 and the independent variables included *Good offers seen* (in the Regret Feedback treatment) or *Good offers rejected* (in the No Feedback treatment). We found no significant effects. See Appendix 4, Section 11. This is further evidence that comparisons between responses to price search and lottery tasks are not distorted by order effects.

they took account of the value of the best previous free-recall offer. The implication is that subjects were using quite sophisticated heuristics that were capable of identifying the main determinants of the distributions of earnings implied by the two options they faced. Even so, if our findings are viewed in the perspective of behavioural economics, in which the context-dependence of preferences is a recurring theme, the individual-level consistency of attitudes to risk between price search and lottery tasks is a striking regularity that needs to be explained. Part of the explanation may be that the risky options in our experiment (i.e. rejecting the time-limited offer or choosing Lottery 2) had relatively low variance. Thus, these tasks did not give much scope for the under- and over-weighting of small probabilities that is a major cause of the description-experience gap (see Section 5.5).

The final unexpected feature of our results is that time-limited offers were more likely to be chosen when the time available for decision-making was longer. Leaving aside the 2 s tasks, where the time constraint was so short that it often cut in while deliberation was in progress, we found no direct evidence that subjects used simpler decision heuristics when time constraints were tighter. Irrespective of the tightness of the time constraint, subjects took account of the same three items of task-specific information (the position and price of the time-limited offer and the value of the best preceding free-recall offer), ignored the same fourth item (the nature of the feedback they would receive), and adapted their decisions in the light of realisations of offer values in previous tasks. The degree of dispersion in responses and the (very low) frequency of choices of dominated offers were very similar across the range of time intervals. These qualitative similarities between decisions made under very different time constraints suggest that, in our experiment, the higher rate of rejection of time-limited offers under tighter time constraints was probably not the result of any fast-and-frugal heuristic, such as the 'run from time-limited offers' rule discussed by Brown et al. (2018), though future research may be useful.

Nevertheless, price search decisions were more risk-averse when the time interval was longer. The most natural interpretation of this finding is that longer time intervals allowed subjects to take account of some additional factor that favoured the choice of the time-limited offer. Our post hoc conjecture is that the special attraction exerted by time-limited offers is related to the distinction between the certainty of the time-limited offer price and the uncertainty about the price that would be paid if that offer was rejected. Two features of the framing of price search problems – not only in our experiment, but also in natural-world retail markets – may make this distinction particularly salient. First, at the time the accept/reject decision has to be made, the certainty has a concrete representation as a visible offer, while the uncertain alternative prices have to be imagined. Second, if the time-limited offer is rejected, discovering the best alternative offer is a process that is extended over time and requires active involvement by the subject or consumer. These features – neither of which is present in the lottery-choice tasks standardly used to elicit attitudes to risk – dramatise the underlying uncertainty of price search problems and may be a cue for negative emotional responses. As a result of these features, deliberation about what is involved in accepting and rejecting the time-limited offer may prime attitudes of risk aversion. The longer a subject has to think about the decision to accept or reject, the more time there is for this priming to operate.

8. Conclusion

In existing industrial organisation models of time-limited offers, retailers use such offers as a means of deterring search and so softening competition. This search deterrence effect is mediated by consumers' binary decisions between, on the one hand, the certain outcome of accepting a time-limited offer and, on the other, the uncertain prospect implied by rejecting that offer and continuing to search. These models assume that, in their search behaviour, consumers are risk-neutral maximisers of expected utility. Our main objective was to investigate whether, other things being equal, 'behavioural' consumers are more or less likely to accept time-limited offers than their expected-utility-maximising counterparts. Our experimental results suggest that they are *more* likely to accept such offers.

However, we found no support for the hypothesis that time-limited offers are attractive to consumers as a way of avoiding regret. Nor did we find evidence that tight time constraints, considered in isolation, make time-limited offers more likely to be chosen. Indeed, we found the opposite effect: holding constant the value of the time-limited offer and the probability distribution of the offers that might be discovered by further search, subjects who had less time to think about their decisions were *less* likely to accept time-limited offers. Even when subjects were allowed only 4 s in which to accept or reject an offer, their decisions were remarkably sophisticated, responding in a qualitatively rational way to the information contained in the offers they had seen.

One should be wary of extrapolating too directly from laboratory behaviour to real markets, but our results suggest that the main mechanism by which time-limited offers benefit retailers is not by using time constraints to impair the rationality of consumers' decision-making. Rather, it is a mechanism of search deterrence – a restriction of consumers' opportunities to compare available offers – amplified by consumers' risk aversion. The evidence of our experiments suggests that problems of choosing between accepting and rejecting time-limited offers are particularly likely to induce risk aversion. Risk aversion is not attenuated, and may even increase, if consumers are able to spend more time thinking about whether to accept a time-limited offer.

Declaration of Competing Interest

None.

Acknowledgements

We thank the editors, two anonymous referees and participants at various presentations for constructive comments, Peter Moffatt for advice on experimentics, and Cameron Belton and Lian Xue for experimental assistance. The research was supported by the [Economic and Social Research Council](#) of the UK (awards no. [ES/K002201/1](#) and [ES/P008976/1](#)) and by the Centre for Behavioural and Experimental Social Science (CBESS) at the [University of East Anglia](#) (UEA). Sugden's work was also supported by funding from the [European Research Council](#) (ERC) under the European Union's [Horizon 2020](#) research and innovation programme, grant agreement no. [670103](#). Ethical approval was obtained from the UEA School of Economics Ethics Committee.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.jebo.2019.09.008](https://doi.org/10.1016/j.jebo.2019.09.008).

References

- Armstrong, M., Zhou, J., 2016. Search deterrence. *Rev. Econ. Stud.* 83, 26–57.
- Bell, D., 1982. Regret in decision making under uncertainty. *Oper. Res.* 30 (5), 961–981.
- Ben Zur, H., Breznitz, S.J., 1981. The effect of time pressure on risky choice behavior. *Acta Psychol. (Amst.)* 47 (2), 89–104.
- Bock, O., Baetge, I., Nicklisch, A., 2014. hroot: Hamburg registration and organization online tool. *Eur. Econ. Rev.* 71, 117–120.
- Bone, J., 2006. *The Hard Sell: An Ethnographic Study of the Direct Selling Industry*. Ashgate, Aldershot, UK.
- Brown, A.L., Viriyavipart, A., Wang, X., 2018. Search deterrence in experimental consumer goods markets. *Eur. Econ. Rev.* 104, 167–184.
- Busemeyer, J., Townsend, J., 1993. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* 100, 432–459.
- Cialdini, R., 2001. *Influence: Science and Practice*. Allyn and Bacon, Needham Heights, MA.
- Diamond, P.A., 1971. A model of price adjustment. *J. Econ. Theo.* 3 (2), 156–168.
- Engelbrecht-Wiggans, R., 1989. The effect of regret on optimal bidding in auctions. *Manage. Sci.* 35, 685–692.
- Finucane, M.L., Alhakami, A., Slovic, P., Johnson, S.M., 2000. The affect heuristic in judgments of risks and benefits. *J. Behav. Deci. Mak.* 13 (1), 1–17.
- Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* 10 (2), 171–178.
- Gigerenzer, G., Goldstein, D., 1996. Reasoning the fast and frugal way: models of bounded rationality. *Psychol. Rev.* 103, 650–669.
- Gilbert, D., Morewedge, C., Risen, J., Wilson, T., 2004. Looking forward to looking backward: the misprediction of regret. *Psychol. Sci.* 15, 346–350.
- Heidhues, P., Kőszegi, B., 2014. Regular prices and sales. *Theo. Econ.* 9, 217–251.
- Hertwig, R., Barron, G., Weber, E., Erev, I., 2004. Decisions from experience and the effect of rare events in risky choice. *Psychol. Sci.* 15, 534–539.
- Huck, S., Wallace, B., 2015. The impact of price frames on consumer decision making: experimental evidence. WZB Berlin Social Science Center. Earlier version published as Report OFT 1226 Working paper. Office of Fair Trading.
- Humphrey, S., 2004. Feedback-conditional regret theory and testing regret-aversion in risky choice. *J. Econ. Psychol.* 25 (6), 839–857.
- Ibanez, M., Czermak, S., Sutter, M., 2009. Searching for a better deal: on the influence of group decision making, time pressure and gender on search behaviour. *J. Econ. Psychol.* 30, 1–10.
- Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47, 263–291.
- Kocher, M.G., Pahlke, J., Trautmann, S.T., 2013. Tempus fugit: time pressure in risky decisions. *Manage. Sci.* 59 (10), 2380–2391.
- Loomes, G., Moffatt, P., Sugden, R., 2002. A microeconomic test of alternative stochastic theories of risky choice. *J. Risk Uncertain.* 24, 103–130.
- Lippman, S., Mamer, J., 2019. Exploding offers. *Decision Analysis* 9, 6–21.
- Niederle, M., Roth, A., 2009. Market culture: how rules governing exploding offers affect market performance. *Am. Econ. J. Microecon.* 1 (2), 199–219.
- Nursimulu, A., Bossaerts, P., 2014. Risk and reward preferences under time pressure. *Rev. Finance* 18, 999–1022.
- O'Donoghue, T., Rabin, M., 2001. Choice and procrastination. *Q. J. Econ.* 116, 121–160.
- Office of Fair Trading (2010). *Advertising of Prices*. Report OFT 1291. London: Office of Fair Trading.
- Payne, J.W., Bettman, J.R., Johnson, E.J., 1988. Adaptive strategy selection in decision making. *J. Exp. Psychol. Learn. Memory Cognit.* 14 (3), 534–552.
- Rabin, M., 2000. Risk aversion and expected-utility theory: A calibration theorem. *Econometrica* 68 (5), 1281–1292.
- Rapoport, A., Tversky, A., 1970. Choice behavior in an optional stopping task. *Organ. Behav. Hum. Perform.* 5 (2), 105–120.
- Richard, R., van der Pligt, J., de Vries, N., 1996. Anticipated regret and time perspective: Changing sexual risk-taking behavior. *J. Behav. Deci. Mak.* 9 (3), 185–199.
- Roth, A., Xing, X., 1994. Jumping the gun: Imperfections and institutions related to the timing of market transactions. *Am. Econ. Rev.* 84 (4), 992–1044.
- Saqib, N., Chan, E., 2015. Time pressure reverses risk preferences. *Organ. Behav. Hum. Decis. Process.* 130, 58–68.
- Salop, S., Stiglitz, J., 1977. Bargains and ripoffs: A model of monopolistically competitive price dispersion. *Rev. Econ. Stud.* 44, 493–510.
- Schunk, D., Winter, J., 2009. The relationship between risk attitudes and heuristics in search tasks: a laboratory experiment. *J. Econ. Behav. Org.* 71 (2), 347–360.
- Seale, D., Rapoport, A., 1997. Sequential decision making with relative ranks: an experimental investigation of the 'secretary problem'. *Org. Behav. Hum. Decis. Process.* 69 (3), 221–236.
- Simon, H., 1957. *Models of Man, Social and Rational*. Wiley, New York.
- Simonson, I., 1992. The influence of anticipated regret and responsibility on purchase decisions. *J. Consum. Res.* 19 (1), 105–118.
- Sitizia, S., Sugden, R., 2011. Implementing theoretical models in the laboratory, and what this can and cannot achieve. *J. Econ. Methodol.* 18 (4), 323–343.
- Sonnemans, J., 1998. Strategies of search. *J. Econ. Behav. Org.* 35 (3), 309–332.
- Spybrook, J., H. Bloom, R. Congdon, C. Hill, A. Martinez, and S. Raudenbush (2011). *Optimal Design Plus Empirical Evidence: Documentation for the 'Optimal Design' Software Version 3.0*. Available from www.wtgrantfoundation.org.
- Stewart, N., Chater, N., Brown, G., 2006. Decision by sampling. *Cognit. Psychol.* 53, 1–26.
- Varian, H., 1980. A model of sales. *Am. Econ. Rev.* 70, 651–659.
- Zeelenberg, M., Beattie, J., Van der Pligt, J., de Vries, N., 1996. Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organ. Behav. Hum. Decis. Process.* 65 (2), 148–158.
- Zeelenberg, M., Pieters, R., 2004. Consequences of regret aversion in real life: The case of the Dutch postcode lottery. *Organ. Behav. Hum. Decis. Process.* 93 (2), 155–168.
- Zizzo, D.J., 2005. Transfer of knowledge and the similarity function in economic decision-making. In: Zizzo, D.J. (Ed.), *Transfer of Knowledge in Economic Decision-Making*. Palgrave Macmillan, London, pp. 1–27.