

CSAE Working Paper WPS/2025-06

Ignorance is bliss? Rejection and discouragement in on-the-job search

Rocco Zizzamia

University of Oxford

rocco.zizzamia@economics.ox.ac.uk

May 21, 2025

Abstract

Searching for jobs often involves repeated rejection. If discouraged searchers reduce search effort in response, this decreases their probability of finding a (good) match, with negative implications for the individual searcher and for the efficiency with which talent is allocated to jobs in general. Using a lab-in-the-field experiment with young workers in South Africa, I examine whether experiencing repeated rejection discourages further search. Participants repeatedly choose between two tasks: a high-return task with frequent feedback containing rejection signals, and a low-return task without immediate rejection feedback. By experimentally varying monetary rewards and rejection exposure, while controlling for learning and risk preferences, I isolate the psychological cost of rejection as a driver of search behaviour. I find that subjects choose to reduce their expected earnings to avoid rejection signals. This behaviour suggests that rejection imposes a psychological cost that motivates active information avoidance and decreases job search.

JEL: D91, J63, J64

Acknowledgements: I am very grateful to Christopher Adam, Kai Barron, Ihsaan Bassier, Joshua Budlender, Gauri Chandra, Stefan Dercon, Nilmini Herath, Vatsal Khandelwal, Lukas Lehner, Craig McIntosh, Muhammad Meki, Sandra Sequeira, Christian Meyer, Kate Orkin, Ragan Petrie, Patrizio Piraino, and participants at the Jobs and Development Conference, the Oxford CSAE Workshop, Oxford Firms and Labour Workshop and the ESA Africa Conference for helpful comments and discussions. I am especially thankful to Pramila Krishnan. This research was possible thanks to partnerships with CapaCiTi, the Youth Employment Service, and Younglings Africa, with special thanks to Fiona Tabraham. I gratefully acknowledge funding for this research from the Center for the Study of African Economies, Oxford. Any errors are my own.

1 Introduction

Being rejected hurts. Activities like job search, applying for graduate schools, or dating inherently involve exposure to frequent rejection, yet offer potentially high rewards—finding a good job, attending a top school, or meeting a suitable partner. In the labour market context, psychological barriers may significantly affect job seekers’ behaviour, particularly their willingness to continue searching after repeated rejections. Although existing research has examined various barriers to job search, the role of psychological barriers—specifically rejection-induced psychological costs—remains understudied (Caria et al., 2024). Early evidence suggests that interventions reducing psychological costs increase job search effort and effectiveness (Field et al., 2023). However, little empirical research directly examines how repeated rejection itself shapes job search behaviour.

This paper investigates whether job seekers deliberately reduce their search effort—and thus forgo potential earnings—to avoid repeated exposure to rejection signals, even when monetary incentives favour continued search. Conceptually, repeated rejection may be especially discouraging in slack labour markets, where the job application success rate is lowest. Receiving repeated rejection signals might threaten self-esteem or self-perception of employability, creating psychological incentives to avoid further rejection by limiting search effort.

To empirically isolate the psychological cost of rejection, I conduct a lab-in-the-field experiment with young South African workers in Cape Town, generating thousands of search decisions from 203 subjects using a within-subjects design. Subjects play a multi-round game in which they repeatedly choose between two tasks: ‘searching for a new job’ (henceforth ‘search’) or ‘focusing on the current job’ (henceforth ‘stay’, with identical probabilities of payoff conditional on effort. The ‘search’ task generates frequent rejection signals while the ‘stay’ task generates infrequent rejection signals. In the benchmark experimental condition the payoff is 10% higher in the ‘search’ task than it is in the ‘stay’ task. While the absolute difference in expected returns between tasks is small, the average subject is induced to search by these monetary incentives. Crucially, subjects receive complete information upfront on payoffs and success probabilities, eliminating the role of learning through rejection as a confounder.

The critical distinction between tasks is the timing and frequency of feedback about rejection. In the ‘search’ task, subjects receive round-by-round rejection feedback; in

the ‘stay’ task, rejection feedback is delayed until the end. Thus, choosing ‘search’ involves frequent rejection signals, whereas choosing ‘stay’ minimises immediate exposure to rejection. Across experimental conditions, I vary monetary rewards and rejection exposure levels, enabling measurement of how subjects trade-off expected earnings against avoiding rejection signals. This design allows me to observe active information avoidance behaviour (Golman et al., 2017), interpreted as a willingness to pay to avoid receiving negative information.

In the first round of the game, prior to receiving any rejection signals, subjects respond strongly to monetary incentives favouring search, with 65% opting to search. However, being exposed to frequent rejection signals, subjects systematically switch from the ‘search’ task to the lower return ‘stay’ task. Within six rounds of the game, search rates decline to 45%. In other words, the average subject becomes willing to sacrifice expected earnings to minimise exposure to rejection signals. On average, search rates are between 21-26 percent lower when subjects are exposed to rejection compared to experimental conditions in which rejection exposure is removed.

By varying rejection frequency independently of task framing, I am able to rule out the hypothesis that framing effects may explain these results. Additionally, by providing detailed upfront information and through regular comprehension checks, my design overcomes concerns that mistaken beliefs about returns to search may be driving results.

Conceptually, I interpret rejection avoidance as active information avoidance (Golman et al., 2017; Andries and Haddad, 2020). Repeated rejection may threaten valued self-beliefs, prompting individuals to strategically avoid further negative signals. This interpretation aligns with behavioural theories positing costly information avoidance to protect self-esteem (Dana et al., 2007; Andreoni et al., 2017; Momsen and Ohndorf, 2020; Kőszegi et al., 2021). I lay out a simple theoretical framework which explains the paper’s empirical results in this light.

This paper contributes to literature on psychological barriers in job search, particularly discouragement effects (Rosholm and Toomet, 2005; Bjørnstad, 2006; Burns et al., 2010; Wheeler et al., 2022), and recent work examining inadequate on-the-job search (Spinnewijn, 2015; Abel et al., 2019; Jäger et al., 2021; Field et al., 2023). My findings highlight repeated rejection as a significant psychological barrier to sustained job search, particularly relevant in slack labour markets characterized by high rejection rates (Kerr, 2018; Zizzamia and Ranchhod, 2019; Donovan et al., 2020). To the extent that

discouragement in job search has been studied in economics, this has tended to focus on discouragement among *unemployed* workers (Rosholm and Toomet, 2005; Bjørnstad, 2006; Burns et al., 2010; Berg and Uhlendorff, 2018; Mueller et al., 2021). However, discouragement may occur also for those who are searching on-the-job, with implications for employment-to-(un)employment transitions rather than unemployment duration. Reduced on-the-job search may also have implications for monopsony power Jäger et al. (2021): If workers under-invest in exploring their outside options, the terms and wages that employers can offer to keep workers on may be lower than they otherwise might be.

This paper also contributes to a large and growing literature on active information avoidance in economics and behavioural science (see Golman et al. (2017) for an review). Research on the topic has shown that people are willing to incur a cost to avoid receiving information that would threaten a valued aspect of their self-image or damage their self-esteem (Dana et al., 2007; Andreoni et al., 2017; Andries and Haddad, 2020; Momsen and Ohndorf, 2020; Kőszegi et al., 2021). The behavioural game I implement can be framed as an investigation into active and costly information avoidance, where rejection is merely a form of information, and choices *not* to search are costly in terms of expected earnings. While Andolfatto et al. (2009) and Kőszegi et al. (2021) have developed theoretical models which suggest that information avoidance may disincentivise job search, to my knowledge this paper is the first to experimentally test this prediction. In addition, current work on willful information avoidance largely approaches the issue as a static phenomenon, assuming that people already know their preferences for seeking or avoiding information (Dana et al. (2007) and Andreoni et al. (2017), inter alia). I add to this literature by approaching willful and costly information avoidance as a dynamic phenomenon in which information avoidant behaviour evolves as people, through exposure to information, learn about and/or change their preferences for acquiring or avoiding information.

The paper proceeds as follows: Section 2 provides empirical and theoretical context. Section 3 describes the protocol of the laboratory experiment in some detail. Section 4 describes the four treatment conditions. Section 5 presents the main results, starting with a descriptive sketch of search behaviour in the laboratory and proceeding to the causal estimation of exposure to rejection signals on search rates. In Section 6 I propose a simple theoretical framework through which I interpret the results of the game. In Section 7 I discuss the main results in the context of the theoretical framework I introduce in the previous section, and discuss the potential role of confounding behavioural explanations for the results I observe. Section 7 concludes.

2 Context

2.1 The South African labour market

South Africa is experiencing a protracted unemployment crisis. As of the second quarter of 2022, the unemployment rate was estimated to be 35,3 percent – the highest since Statistics South Africa began collecting quarterly labour force survey data in 2008 (Statistics South Africa, 2022). If the expanded definition of unemployment is used, which includes ‘discouraged work-seekers’, the rate is 46,2 percent. Unemployment disproportionately affects Black South Africans (39,1 percent), women (38,2 percent) and youth, (66,5 percent for those between 18-24, and 43,5 percent for those between 25-34).

Using panel data, researchers estimate that between a quarter and a third of South Africans are stuck in chronic unemployment, with little chance of finding stable employment (Ingle and Mlatsheni, 2017; Zizzamia and Ranchhod, 2019). At the same time, South Africa has a very high rate of employment volatility, with many among the employed losing their jobs over time, and some among the unemployed moving into employment, even if temporarily (Banerjee et al., 2008; Kerr, 2018). Young workers’ position in the labour market is especially precarious: Jobless and hence unable develop post-schooling human capital through work-based-learning, and struggling to signal the human capital they already possess to prospective employers, young workers are vulnerable to becoming trapped in a state of chronic joblessness.

It is in the context of these challenges that fixed-term employment opportunities, including internship-style programmes, hold some promise. Simply getting a foot in the door may substantially alleviate several constraints, through the work-based learning led accumulation of human capital, networking effects and the ability to provide positive signals to future employers. These short-term opportunities are stepping-stones, rather than the final destination. In order to sustain an upward career trajectory, workers need to proactively seek employment opportunities for when their current contract ends. While some workers may climb the job ladder at their current place of work—perhaps receiving a contract extension or the offer of a permanent position—many others will need to find new jobs, which will require on-the-job search if workers aim to avoid falling into unemployment.

It is helpful to know something of the origin story of this paper: The idea emerged from

a previous research project in which I was involved in designing an experimental evaluation of an internship programme for unemployed youth in South Africa. In this programme, internships were offered as 12 month contracts, with no guarantee of retention. In the course of designing the evaluation, we implemented a structured survey with participating interns, and did qualitative interviews with both participating firms and interns. I was struck by three findings: First, interns were optimistic about their chances of being retained at the firm providing the internship—over 50% of surveyed interns considered it highly likely they would be offered a permanent position. Second, the true retention rates for interns in their host firms was less than 5%, suggesting dramatic over-optimism among interns. Third, from focus group discussions and from conversations with firm hiring managers, it appeared that on-the-job search among interns was very low. Taken together, these three findings may have helped explain the fact that approximately 50% of interns fell into unemployment after completing the internship. This puzzle prompted me to consider the barriers to on-the-job search, one aspect of which this paper addresses.

2.2 Rejection avoidance as information avoidance

The rejection avoidant behaviour I study in this paper ought to be understood as an example of ‘active information avoidance’. Golman et al. (2017) define active information avoidance as those cases in which individuals fail to acquire information even when *a*) they are aware that the information is available, and *b*) the individual has free access to the information. Active information avoidance has been well studied in various domains in economics and behavioural science, including in stock market investors displaying an ‘ostrich effect’ when markets are down (Karlsson et al., 2009), individuals avoiding potentially useful medical testing (Oster et al., 2013; Ganguly and Tasoff, 2017), consumers avoiding information on the environmental impact of their choices (Momsen and Ohndorf, 2020), and altruistically-minded people being willing to incur costs to avoid being asked to donate to charity, thereby ‘[avoiding] empathic stimulation to regulate their giving and guilt’ (Andreoni et al. (2017) p. 625). Dana et al. (2007) show that in a dictator game, subjects prefer to remain ignorant of the effects of their choices on another subject.

One explanation for this information avoidance is that people strategically avoid being exposed to information where receiving this information would force them into a trade-off between taking costly action or revising a belief from which they derive positive utility. Both decisions incur a cost – in the one case monetary, and in the other psychological.

Relevant beliefs may include those which are relevant to one’s self-concept, such as self-esteem.¹

Rejection in job search, like receiving an unfortunate medical diagnosis, is information which is both potentially useful, and which is unpleasant to receive. Rejection may contain instrumentally useful signals – for instance, on comparative advantage, match quality, and returns to certain search strategies – which may improve the efficacy of job search. However, the information contained in rejection may also be dissonant with beliefs from which individuals derive positive utility: for instance, these might include beliefs that individuals have of themselves as worthy, competent and in control of the outcomes they aspire to. To the extent that individuals derive utility from these beliefs, in a way that is analogous to medical screening, it is an empirical question whether, through rejection, the hedonic disutility of a possible downward revision is greater than the utility of the potential instrumental value of the information. If the anticipated disutility of a downward revision of beliefs is great enough, we might expect individuals to strategically avoid exposure to rejection, even if this information avoidance comes at a material cost.²

3 Empirical approach and data

3.1 Sample

The experiment was conducted in Cape Town, South Africa, in August and September 2022. Experimental sessions were hosted in-person in computer laboratories on the campus of the implementing partner, a digital skills training center which hosts young South Africans from low socio-economic status backgrounds on three to six month paid ‘learnerships’.

I deliberately selected this sample for its relevance to the research question: The

¹Similar behavioural distortions aimed at maintaining a valued self-image have also been shown to occur ex post (Kőszegi et al., 2021). For instance, people have been shown to distort memories which are dissonant with a valued view they have of themselves (Huffman et al., 2022).

²Repeated rejection may also make the information on rejection more *salient* than the explicit goal of the game – to maximise earnings. The salience of the rejection stimulus may interfere with the income-maximising goal, leading to suboptimal decision-making (Bordalo et al., 2022). It is possible that salience may be at play in this context, but, as I argue later, is not incompatible with the information-avoidance story.

subjects I worked with were mostly employed on a short-term contract and faced uncertain prospects in terms of finding a job or being retained or ‘placed’ after their learnership. They were therefore grappling, in their own lives, with the problem being played in the game: how to allocate effort between focusing on their current job in the hope of increasing their chances of being retained or placed, or to searching for a new job.³ In total 203 subjects participated in the experiment. Of these subjects, 158 were drawn from the learnership programme at CapaCiTi and were thus in a form of temporary employment. The remainder were unemployed, and were drawn from CapaCiTi’s alumni database or from the networks of the employed subjects.

Table 3.1 presents basic descriptive statistics on the laboratory sample. Table 3.1 also puts the sample in the broader context of the South African labour market by including averages for a relevantly similar sample (column 2) and for national adult averages (column 3), using 2017 data from the National Income Dynamics Study (NIDS) – a large nationally representative survey.

The CapaCiTi sample is young and predominantly black African. Women make up slightly more than half of the sample. Most participants have some prior work experience and have undergone some form of post-school training. This sample is typical of a particular class of urban youth: educated enough to aspire to stable white-collar work, but struggling to develop and signal their human capital. As a result, they cycle between states of precarious short-term employment and unemployment.

Column 2, which is restricted to a NIDS sample of urban non-white adults under the age of 28, puts the lab sample into perspective. In terms of demographics, the lab sample is slightly less black African (likely due to its location in Cape Town, which has a larger coloured population⁴) and substantially more educated, with approximately double the rate of post-school training. In the labour market, the two samples are fairly distinct, reflecting the selection of the lab sample through a learnership programme. This explains the higher employment rate (78 percent compared to 45 percent) and substantially lower earnings conditional on employment (R3,436 compared to R6,140). For completeness, Column 3 reports averages for a nationally representative sample of the South African working-age population (ages 18–60).

³It is worth noting that, perhaps for this reason, subjects had little trouble understanding the set-up of the game or the choices being presented.

⁴Coloured is official South African terminology for an ethnic group of mixed racial origin.

Taking a dynamic view, the ‘lab comparable’ sample exhibits higher rates of chronic unemployment (47 percent of unemployed individuals were unemployed in both 2014 and 2017, compared to 35 percent in the full sample) and lower rates of stable employment (25 percent compared to 40 percent). However, as expected for a younger sample, rates of job finding—likely reflecting labour market entry—are higher (20 percent versus 14 percent).

Table 3.1: Lab sample compared to National Income Dynamics Study samples

	Lab sample	NIDS sample (young, urban)	NIDS sample (all adults)
<i>Demographics</i>			
Black African	0.81 (0.39)	0.87 (0.33)	0.88 (0.33)
Woman	0.59 (0.49)	0.50 (0.50)	0.50 (0.50)
Age	23.42 (2.69)	23.90 (3.38)	34.49 (11.13)
Has some tertiary education	0.67 (0.47)	0.17 (0.37)	0.17 (0.37)
<i>Labour market</i>			
Employed	0.78 (0.42)	0.45 (0.50)	0.54 (0.50)
Not economically active	0.14 (0.35)	0.36 (0.48)	0.31 (0.46)
Unemployed	0.08 (0.27)	0.19 (0.39)	0.15 (0.36)
Real monthly earnings (ZAR)	3,436.57 (971.43)	6,140.22 (5,280.22)	7,235.30 (9,376.40)
Past work experience (months)	10.86 (12.88)	-	-
<i>Household variables</i>			
Real hh monthly earnings (ZAR)	-	9,504.75 (11402.83)	9,772.30 (13598.18)
Number of household residents	-	4.17 (2.81)	4.39 (3.23)
Real p.c. hh monthly income (ZAR)	-	3,755.96 (5,663.18)	3,712.09 (6,407.40)
Real p.c. hh expenditure (ZAR)	-	2,872.24 (7,715.30)	2,550.46 (4,080.49)
<i>Employment dynamics: 2014-2017</i>			
Lost job between 2014-17	-	0.08 (0.28)	0.11 (0.32)
Employed in 2014 and 2017	-	0.25 (0.43)	0.40 (0.49)
Unemployed in 2014 and 2017	-	0.47 (0.50)	0.35 (0.48)
Found job between 2014-17	-	0.20 (0.40)	0.14 (0.35)
Observations	203	4484	16036

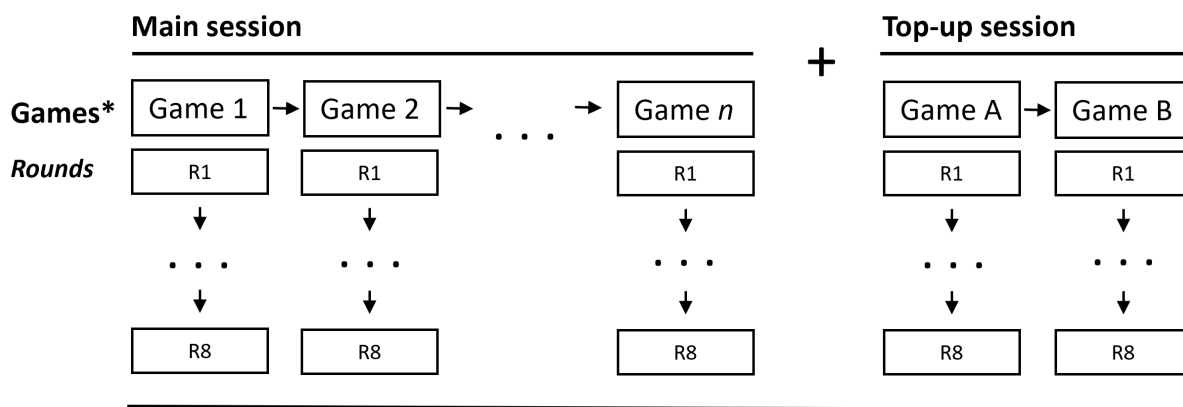
Notes: Standard deviations are reported in parentheses. The first column reports averages and standard deviations for the laboratory sample. Columns 2 and 3 report statistics from the 2017 Wave of the National Income Dynamics Study. For comparability with the lab sample, the sample in Column 2 is restricted to urban non-white adults under the age of 28. Statistics reported in columns 2 and 3 are weighted using NIDS post-stratified weights. All monetary figures are deflated to August 2022.

3.2 The structure of the experimental sessions

The experimental sessions have a simple hierarchical structure, represented graphically in Figure 1: Games – representing treatment conditions described below – are nested within the overarching session, and rounds are nested within games. The overarching experimental session lasted on average 90 to 120 minutes. Each subject was eligible to play each game only once. Subjects were invited to complete two unsupervised top-up games, which they could access remotely, but only if they had already completed the main session in person. Details on this top-up session are provided in Appendix A. Henceforth, my definition of ‘experimental session’ will include top-up sessions for those who completed them.

Each experimental session was made up of between five and six games, representing the treatment conditions described in Table 4.1 and Appendix B. In order to rule out possible sequencing effects, the order in which the games were played was randomised in each experimental session. A total of 10 sessions were conducted.

Figure 1: Session structure



*Game order is randomized across sessions

Each experimental session involved between 7 and 35 subjects. All decisions were made on a computer screen. Participants were seated with an empty seat separating one another so as to maintain privacy and ensure they could not observe each others’ decisions or the outcomes of their games. Verbal communication of decisions and outcomes was not permitted.

At the beginning of the experimental session, subjects were guided through a set of

pages which described the game and provided a comprehensive set of instructions. I personally guided all the experimental sessions. The introduction also included a shortened practice round in which subjects were guided through the choices in the game, and could practice the real effort task (described in subsection 3.4 below). Subjects also had the opportunity to ask questions during this introductory session.

Apart from a practice round⁵, all decisions had monetary consequences. In each game, subjects were eligible to receive a pay-off of either ZAR55 (US\$3.30) for the higher paying option, ZAR50 (US\$3.00) for the lower paying option, or nothing.⁶ Subjects were made aware that the probability of receiving a payoff was a function of the performance on an effort task, and luck. Since each session consisted of multiple games, the subject's earnings over the session were simply the sum of earnings across all games, plus the fixed cash show-up fee of ZAR80 (US\$4.50)⁷ This implies a potential of earning at most ZAR380 (US\$23.00). The average subject received approximately ZAR160 (US\$9.00) per session, made up of ZAR80 (US\$4.50) as the show-up fee, and ZAR80 (US\$4.50) in winnings.⁸

Following the introduction, practice round and comprehension test, subjects could begin playing the games. In most sessions subjects completed five or six separate games. After each game, subjects were informed of their earnings for that particular game, and after the final game they were informed of their earnings for the full session. Total earnings is simply the sum of earnings across all games played.

After completing the games, subjects were asked to complete a survey collecting information on demographic characteristics, labour market history and behaviour, risk and time preferences, and psycho-social characteristics. After having completed the survey they received their compensation and pay-off and could leave.

The behavioural game and accompanying survey were designed and implemented in

⁵Practice data is excluded from the study.

⁶The difference in monetary incentives between tasks is small. This difference was intended to strike a balance between two considerations: On the one hand, a difference that was too large would pose the risk of dominating all other motivations. On the other hand, the difference would need to be large enough to induce a preference for one task over the other. During piloting I determined that a difference of 10% was large enough to induce subjects to prefer the higher-return task initially, but small enough that subjects might forego this premium based on exposure to rejection during the game.

⁷The relatively large show-up fee was chosen for fairness concerns – the experiment's design was such that I expected a substantial minority of subjects not to receive any payoff through the game itself, so guaranteeing an acceptable base payment was deemed to be important by myself and the implementation partner.

⁸Participants were compensated for their transport expenses.

oTree (D. Chen et al., 2016).

3.3 Game protocol

For clarity of exposition, in this subsection I only describe the ‘base game’. As described in the preceding subsection, in each experimental session subjects play several variants of this base game, in each of which a key parameter is varied. These variations are described in the next section.

For a subject, the (stated) objective of the game is to maximise earnings. In each round, subjects choose between two job offers: they can seek employment either by trying to find a new job *by searching* (‘search’), or by trying to secure a contract extension *by focusing on their current job* (‘stay’).⁹ In the base game, ‘searching’ strictly dominates ‘staying’ in terms of expected pecuniary returns. Therefore, in expectation, subjects will achieve the payoff-maximising objective by always searching.

In the base game, this choice is made eight times. Having made the decision, the subject then completes a real effort task (described in subsection 3.4 below). Outcomes – that is, whether or not a subject receives a job offer in their given choice – depends on the subject’s performance in the real effort task, and on luck (described in more detail in subsection 3.4 below). In addition, subjects are informed that their chances of ‘finding a new job’ or ‘securing a contract extension’ are exactly equal for a given level of effort they exert in the real effort task.

Apart from differences in pay-off, the main difference between ‘searching’ and ‘staying’ is that if a subject ‘searches’ she receives a verdict on the outcome of her choice every round (at median effort this is a rejection in 96 percent of cases), compared to if a subject ‘stays’, in which case information on the outcome of her decision is withheld until the end

⁹The choice to frame the scenario explicitly in labour market terms is deliberate: I am seeking to understand labour market behaviour in the first instance, and hence prefer to frame the experiment in these terms rather than use neutral language. The choice of subjects (young entry-level white collar workers) is made for the same reason – for their relevance in exploring this particular labour market dynamic. However, by comparing *SearchNoRej* and *Search* conditions, I also seek to identify whether this explicit framing has an impact on behaviour in the laboratory setting.

of the game (at median effort this is a rejection in 72 percent of cases).¹⁰ The *cumulative* probability of receiving an offer is, however, identical across the two activities. There is no penalty to switching between activities during the game.

Rejection is defined as an unsuccessful attempt at securing a pay-off, in the game framed as a ‘job’. For each subject in each round, I record the decision made (‘search’ or ‘stay’), performance in the real effort task, and the outcome of the round (successful or unsuccessful). When subjects decide to search, the outcome is made visible to them. Unsuccessful outcomes are framed as rejections. Recording this information allows me to observe how decisions and effort in subsequent rounds are affected by information provided on outcomes in prior rounds.

One of the key objectives of the game is to shut down the channel through which subjects learn, by experiencing rejection, about the mechanisms determining outcomes. To do this, I make sure that subjects in both the Search and Stay conditions have access to exactly the same information on the mechanisms determining outcomes: I provide high frequency information on subjects’ performance, the likelihood of success, and the relationship between performance and the likelihood of success in a way that is *identical* regardless of whether a subject chooses to search or not to search. The only thing that differs is the information that is provided on the outcome – which in itself carries no additional information on the underlying probabilities or on the way the game works.

Each round consists of three steps (depicted in Figure 2):

1. **Choice:** Subjects are prompted to make a choice between *searching for a new job*, or *focusing on their current job*. Each option is associated with an identical cumulative probability of receiving a job offer, but a different payoff conditional on receiving an offer. This is communicated clearly to subjects. Participants are tested

¹⁰This round-by-round rejection rate may appear high, potentially raising concerns that if rejection rates in real life are substantially lower, behaviour in the game may display qualitatively different patterns than behaviour in the real world. I selected this high rejection rate for two reasons, the first empirical and the second pragmatic: 1) At least in the South African context, a 96 percent rejection rate per application does not seem unrealistically high. One recent study on job search behaviour of a demographically similar sample in South Africa found that the average jobseeker submitted 10 job applications per week, and received only 0.17 job offers in the last 30 days (Garlick et al., 2023). This implies a success rate per application of only 0.4 percent. Even if we assume that this estimated success rate is a substantial underestimate, the 4 percent success rate in the game does not appear to be unrealistically low. 2) I also chose this low acceptance rate for pragmatic reasons: I effectively ‘lose’ sample members when they are successful when searching for a new job, so a high rejection rate means that I am able to retain more of the sample.

on their comprehension of this concept before starting each game.

2. **Effort task:** Regardless of what subjects choose, they complete a real effort task which determines their likelihood of receiving an offer. The effort task is a simple slider activity based on Gill and Prowse (2019), where subjects have 40 seconds to complete a maximum of 30 sliders. This task was calibrated through piloting so that on average subjects would complete 12 sliders, yielding a round-by-round job offer rate of 4 percent. Probabilities of receiving an offer were transparently determined by performance on the real effort task – each slider represented a 1/3 percentage point increase in the probability of receiving an offer, whether the decision was made to search or not to search. Comprehension checks ensured subjects understood this relationship.
3. **Results:** Subjects are shown a results screen.
 - (a) If they had chosen *not to search*, they are informed only how many sliders they completed and the corresponding probability of receiving an offer. Information on whether or not they received an offer is withheld until the end of the game.
 - (b) If they had chosen *to search*, they are informed whether or not they received an offer, in addition to being told how many sliders they completed and the corresponding probability of receiving an offer.¹¹

¹¹The exception to this is the *SearchNoRej* condition, in which respondents are not informed in this step whether or not they received an offer. See Table 4.1.

Figure 2: Anatomy of a round

(a) Decision screen

Do you want to...?

Search for a new job

- You will be **told after each round** whether you found a job.
- You get **R55 airtime** for finding a job.

Focus on current job

- You will be **told at the end of the entire game** whether or not you get offered an extension to your current job.
- You get **R50 airtime** for keeping your job.

In both cases: More sliders = better chances!

Completing 3 sliders increases your chances of finding a job / keeping your job by 1/100 (from zero)! So...

3 complete sliders = 1/100 chance of being offered a job.

30 complete sliders = 10/100 chance of being offered a job.

What will you choose?


- Search for a new job
- Focus on current job, don't search

Next

(b) Real effort task

Round 1: You are searching for a new job

Time left to complete this page: 0:21



(c) Results screen

Round 1 results: Searching for new job.

You were not offered a job.

You completed **20 sliders**. This means you had a **6/100** chance of being offered a job this round.

Next

For the purposes of the study, once searching subjects receive an offer, I effectively lose the chance to collect data from these individuals for the remainder of that particular game, since they are not longer prompted to make decisions and so I no longer observe the key outcome of interest for these individuals.¹²

Before each game, subjects are also informed of the *game specific characteristics* – described in detail in Section 4 – for instance, the expected pay-off to securing a new job versus keeping their current job, whether or not they will be informed immediately of the outcome of their job search if they decide to search (i.e. exposed to potentially frequent rejection), and over how many rounds the game is played. Because each game differs in these important respects from all other games, before beginning each game a comprehension test was used to ensure that each subject understood the distinctive characteristics of each game. Importantly, subjects were required to correctly identify in which of the two tasks ‘one would earn more’. Subjects could not start the game until they had answered all questions correctly. See Appendix C for a full list of comprehension test questions.

3.4 The real effort task

It is important that subjects perceive outcomes as being determined not only by luck but also by a variable over which they have some control and which is linked to performance or ability. This is because the way that an unsuccessful outcome is interpreted in a pure lottery may be psychologically very different from how an unsuccessful outcome is interpreted where performance is implicated. The real effort task allows for this. In addition, the relationship between effort and outcomes also helps subjects perceive the game as being a plausible simulation of job search in the real world: the real effort task is explicitly framed as representing the effort a job seeker places on looking for work, or that a worker might place in ‘going above and beyond’ to increase their chances of being retained.

¹²If subjects choose to search and are successful in receiving a job offer in a given round, they are no longer prompted to make decisions between ‘searching’ and ‘staying’. Rather, they are informed that they will retain their offer if they maintain a modestly high standard of performance in the effort task going forward. Specifically, subjects keep their offer if they complete as many sliders as in the round in which they were made an offer, minus four. For instance, if a subject receives an offer having completed 14 sliders, she will need to complete 10 in subsequent rounds to keep her offer. In the absence of this minimum performance standard, subjects might be incentivised to search so as to be able to ‘slack off’ once they receive an offer.

Importantly, the relationship between effort and the probability of receiving an offer is transparently revealed to subjects: Specifically, subjects are taught that the probability of receiving an offer is exactly equal at a given level of effort for the two activities of ‘searching’ or ‘staying’. The assumption that subjects understand this feature of the game is important for the causal identification of a rejection effect. It is also important for the effort task to be identical for both ‘search’ and ‘stay’ tasks in order that preferences over the real effort task itself are not driving results.

In behavioural economics and experimental psychology many approaches to measuring effort have been used, including, among others, solving puzzles, doing data entry tasks and cracking walnuts (see Charness and Kuhn (2011) and Charness, Gneezy, et al. (2018) for reviews). I use the Gill and Prowse (2019) ‘slider task’.¹³ In this task, subjects are instructed to move a slider on a computer screen to a specific target point. The screen contains multiple sliders, and the subject’s objective is to correctly ‘complete’ as many sliders as possible within an allotted time, where ‘completing’ means aligning the slider with the target point. A screenshot of the slider task is provided in Figure 2 above, showing that 2/30 sliders have been completed with 21 seconds remaining.

Subjects had 40 seconds to complete up to 30 slider bars, determined through piloting to yield a median slider completion of 12, corresponding to the target 4 percent per-round offer rate.¹⁴ The difference in slider completion rates across activities in the game (i.e. if the subject chose to ‘Search’ or to ‘Stay’) is not significant.

¹³The slider task was first used by Gill and Prowse (2012), and has since been used in a large number of behavioural experiments spanning a broad range of topics (see Gill and Prowse (2019) for a list). The slider task has several advantages which make it an appropriate choice for this experiment: The task is very easy to understand and requires little explanation. It is also able to generate a finely granular measure of effort within a short time frame, allowing for repeated observations for the same individual. It is also easy to implement in a game designed using OTree, and key parameters (such as the number of slides and the allotted time) can be finely calibrated. Importantly, there is little scope for ‘learning’ in the task, meaning that the cost of effort remains roughly constant within individuals over time (Charness and Kuhn, 2011).

¹⁴The probability of receiving an offer in any given round is determined by the number of sliders completed and random chance. Each slider represents a 0.33 percentage point increase in the likelihood of receiving an offer from a base of zero. Therefore, any subjects’ probability of receiving an offer varied from 0 percent (if she completed no sliders) to 10 percent (if she completed 30 sliders). Thus, since the median slider completion rate was 12, the median subject had a per round offer probability of 4 percent.

4 Treatments

The three main treatments of the game create variation in two key parameters, which I use to isolate the effect of rejection on search behaviour. The first source of variation is in the *expected monetary returns to searching for a job*. The second source of variation is in the *amount of rejection that subjects are exposed to when searching for a job*. Table 4.1 summarises the three core treatment conditions.¹⁵ Note that the design and results of a fourth treatment, in which the frequency of rejection was doubled, are discussed in Appendix D.

Table 4.1: Primary treatments

Treatment	Payoff: finding job (ZAR)	Payoff: keeping job (ZAR)	Rejection signals when searching?
1. Stay	50	55	YES
2. Search	55	50	YES
3. SearchNoRej	55	50	NO

Notes: The probability of receiving an offer in any given round is determined by the player’s performance in the effort task in that round, and ranges from 0% to 10% in the 8-round games. At median effort, the *cumulative* probability of receiving an offer over the course of the game is equal in Treatments 1-3 (27.9%). T1 and T2 were pre-specified, while T3 was added during fieldwork. See Appendix B for a discussion.

In *Stay*, there is a monetary incentive to ‘staying’. As long as rejection signals are a net psychological cost, then both rejection effects and cash incentives pull in the same direction. Therefore, both rejection indifferent and rejection averse subjects would prefer to ‘stay’. Rejection aversion as an expression of active information avoidance should be inconsequential in this condition: Expected-payoff maximising subjects should never experience rejection because they should never choose to ‘search’.

In *Search*, which is the ‘base’ game, there is a monetary incentive to ‘search’. If information on rejection enters the individual’s utility function negatively, then rejection information and cash incentives pull in opposite directions. This yields two possible outcomes: If an individual’s aversion to information on rejection is strong enough, then rejection effects will dominate earnings incentives so that the individual will choose to ‘stay’. If, however, an individual’s rejection aversion is not strong enough for rejection

¹⁵This is not an exhaustive list of the treatment conditions implemented in the experiment. Other treatment conditions created variation in the *content* of the rejection messages. These are summarised in Appendix B.

effects to dominate earnings incentives, then the individual will choose to search. Thus, unlike in the *Stay* condition, now information avoidance becomes consequential to search behaviour.

If an individual's incentives to avoid information respond to past exposure to information, rejection effects may start out being dominated by pecuniary incentives, and in later periods may come to dominate pecuniary incentives. Because there is a well-defined monetary premium to searching, we can interpret decisions to stop searching as a willingness to forfeit the known monetary premium in order to avoid information on rejection. Because each game consists of multiple rounds, I will be able to investigate the presence of these dynamic information avoidance effects.

In the experimental set-up, rejection frequency is not the only difference between choosing to 'search' and choosing to 'stay'. Framing effects may also matter, and the high rejection-frequency activity ('searching') is explicitly framed as 'searching for a job' while the rejection-scarce activity ('staying') is explicitly framed as 'focusing on your current job'. To the extent that that subjects carry behaviours, intuitions and norms from the real labour market into the game, this labelling may have a confounding influence on decisions to 'search' or to 'stay', *independent of the experience of rejection*.

The *SearchNoRej* treatment provides a way of overcoming this challenge to the causal identification of a rejection effect. *SearchNoRej* is identical to *Search* in terms of incentives and framing, but differs in that if a subject chooses to search, her rejection exposure is identical to the exposure had she chosen to 'stay'. The difference between *SearchNoRej* and *Search* can therefore be interpreted as the causal effect on search of being exposed to some rejection compared to being exposed to no rejection, having experimentally netted out the framing effects which confound a comparison between *Stay* and *Search*.

5 Results

5.1 Descriptive comparisons of search behaviour across treatments

5.1.1 Treatment 1: *Stay*

In this treatment subjects have a monetary incentive to ‘stay’ (see Table 4.1). Subjects are informed that they will receive ZAR55 (US\$3.30) if they keep their current job, and ZAR50(US\$3.00) if they ‘search’ and succeed in finding a new job, and that their chances of receiving an offer for either are equal (holding effort constant).

The red line in Figure 3a plots the evolution of the proportion of decisions made to search for the eight rounds over which the *Stay* game is played. Search rates remain stable at approximately 30%. Even when subjects are incentivised to ‘stay’, they adopt a mixed strategy—switching between ‘searching’ and ‘staying’, so that search rates do not start at or approach zero. These non-zero search rates are driven both by within-subject switching, and between-subject heterogeneity in terms of game strategies—i.e. the presence of ‘always searchers’ and ‘never searchers’. In this sense, we can consider the observed effect of the incentive in the *Stay* condition to be an upper bound, against which we can compare effects in the *Search* condition (where incentives to search may be offset by the experience of rejection).

5.1.2 Treatment 2: *Search*

In this treatment subjects have a monetary incentive to ‘search’ (see Table 4.1). Subjects are informed that they will receive ZAR50 (US\$3.00) if they keep their current job, and ZAR55 (US\$3.30) if they search and succeed in finding a new job, and that their chances of receiving an offer for either are equal (holding effort constant).

The green line in Figure 3a plots the evolution of the proportion of decisions made to search for the eight rounds over which the *Search* game is played, comparing this to the trends in the *Stay* condition (red line). The bottom two flat horizontal lines plot the average search rate over the eight rounds for the *Search* and *Stay* conditions respectively. The maroon flat horizontal line plots what we can assume would be the average search

rate in the *Search* condition *if subjects respond symmetrically to monetary incentives to search and not to stay*. In other words, this line illustrates how frequently rejection-insensitive subjects would search if they responded symmetrically to inverted monetary incentives.

In the *Stay* condition, over the eight rounds subjects chose to search 30 percent of the time (solid red line). Thus, if subjects responded symmetrically to the inversion of the monetary incentives in the *Search* condition, we would expect search rates in the *Search* condition to be 70 percent (solid maroon line). However, we observe the average search rate in the *search* condition to be 50 percent (solid green line). This implies that the monetary incentive to stay ‘bites’ more than the same monetary incentive to search. This asymmetry suggests that there is a disincentive (by construction, rejection and framing effects) to searching that offsets the monetary incentive to search. Search rates in the *Search* condition are approximately 45 percent lower than would be expected under the assumption of a symmetrical response.

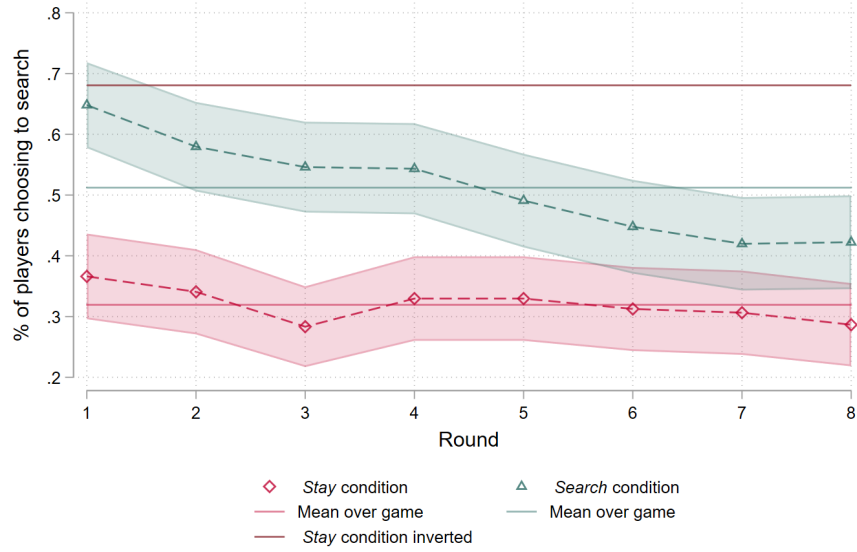
These averages hide important insights which are revealed by an analysis of the dynamics of search rates. Importantly, search rates in the first round in the *Search* condition are exactly what would be implied by symmetry with the *Stay* condition: 65 percent in *Search* round 1 versus 35 percent in *Stay* round 1. In other words, in the first round, the monetary incentive to search appears to ‘bite’ in the same way as the monetary incentive to stay. However, after the first round, there is a sustained decline in search rates in the *Search* condition (dotted green line) while search rates in the *Stay* condition remain relatively stable. Ultimately, by the sixth round the difference in search rates in the *Search* and *Stay* conditions are statistically insignificant. This convergence between the two search trend lines suggests that after repeated rejections, the disincentive to search (plausibly a rejection information avoidance effect) dominates the monetary incentives to search, and induces those subjects initially induced by the monetary incentives to search, to switch strategies.

As discussed in Section 4, rejection aversion may not be the only thing that motivates a subject to decide ‘stay’ – the choices are explicitly framed as ‘searching’ and ‘focusing on your current job’. To the extent that subjects bring norms, habits and intuitions from the labour market to the game, this labour market framing may affect search behaviour independently from the effect of rejection. Figure 3 already provides suggestive evidence that these framing effects are not strong: framing effects are presumably stable over time,

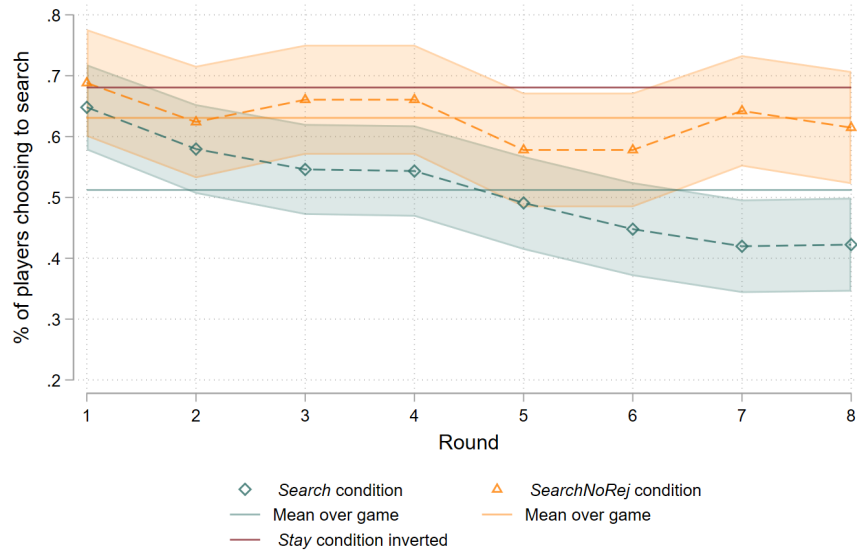
while rejection aversion may change with past experiences of rejection (as discussed in Section 6). The discouragement we observe in the *Search* condition in Figure 3 suggests that the average differences between search rates in *Search* and *Stay* are driven by time-varying factors, not stable factors such as framing effects. Nevertheless, by comparing *Search* and *SearchNoRej*, I provide a cleaner means of identifying the rejection effect.

Figure 3: Search rates

(a) *Stay vs Search*



(b) *Search vs SearchNoRej*



Notes: Shaded areas represent the 95% confidence interval of point estimates in each round. $n = 179\ 109$ for Figures a) and b) respectively.

5.1.3 Treatment 3: *SearchNoRej*

Comparisons between the *SearchNoRej* and *Search* treatments allow us to isolate the effect of rejection on job search. As in *Search*, in this treatment subjects have a monetary incentive *to search* (see Table 4.1). Subjects are informed that they will receive ZAR50 (US\$3.00) if they keep their current job, and ZAR55 (US\$3.30) if they search and succeed in finding a new job, and that their chances of receiving an offer for either are equal (holding effort constant).

The key characteristic of the *SearchNoRej* condition is that, if a subject chooses to search, she does not receive information on the outcome of her decision to search in the Results page after each round (see Figure 2). The outcome on whether or not she received an offer is withheld until the end of the game – just as if she had chosen *not* to search. Since framing is held constant between *Search* and *SearchNoRej*, but the former condition exposes subjects to frequent rejection while the latter does not, a comparison between the search rates in these two conditions yields a clean identification of the rejection effect on search.

Figure 3b compares search rates in the *Search* and *SearchNoRej* conditions. First, we observe that average search rates in the *SearchNoRej* condition (solid orange line) are higher than the search rates in the *Search* condition (solid green line). As we shall see in Table 5.2, these differences are economically and statistically significant. Differences between the average search rates in the *SearchNoRej* condition and the inverted *Stay* condition (solid red line) are smaller and are insignificant.

Search rates for the *Search* and *SearchNoRej* conditions are similar in the first round, but while search rates decline over the eight rounds in the *Search* condition until they fall to approximately 40 percent, in the *SearchNoRej* condition search rates remain stable between approximately 60 and 70 percent.

5.1.4 Comparing subject strategies across games

The preceding analysis gives little insight into the strategies that subjects adopt in the game, and how these strategies respond to the different treatment conditions. Expected-payoff maximising and rejection insensitive subjects would be ‘Always stayers’ in the *Stay* condition, and ‘Always searchers’ in all the other conditions. In reality, there is

heterogeneity both within subjects over time (as subjects start with one strategy and then switch to another during the game, or follow a mixed strategy), and across subjects (with some subjects being more disposed to search than others).

Figure 4 shows the distribution of the proportion of rounds in which subjects chose to search for each of the treatment conditions discussed above. We observe a bunching of search rates around 0 and 100 percent in all four treatments, but with expectedly pronounced bunching at 0 percent in the *Stay* condition, and at 100 percent in the other three conditions. It is also noteworthy that some subjects appear to adopt an inflexible ‘always search’ or ‘never search’ position, persisting in this strategy even in conditions where the monetary incentives pull in the opposite direction. One explanation might be that subjects apply simple heuristics (i.e. ‘always search’ or ‘never search’) for evaluating the decision to ‘search’ / ‘stay’ in each round, which could help explain the bunching at 0 and 100 percent.

Differences between the average search rates across the conditions are therefore driven by those who are induced to switch strategies by the different treatments. The number of ‘never searchers’ is twice as large in the *Stay* condition (red line) as it is in the other three conditions, while this condition has about half as many ‘always searchers’ as the other three conditions. The change in incentives therefore induces some ‘never searchers’ to search, and some ‘sometime searchers’ to become ‘always searchers’.

The difference between the *Search* condition and the *SearchNoRej* condition is slightly more complex: *Search* has a trimodal distribution with a high concentration of ‘sometimes searchers’, while *SearchNoRej* has a bimodal distribution. In other words, there are fewer ‘always searchers’ and more ‘sometimes searchers’ in the *Search* condition relative to the *SearchNoRej* condition. This suggests that fewer subjects who begin searching in the *SearchNoRej* condition are induced to switch strategies during the game than in the *Search* condition – a pattern which is consistent with the discouragement effect of rejection we observe in Figure 3.

Overall, the search patterns in *SearchNoRej* look roughly symmetrical to the search patterns in *Stay* – encouraging, given that fact that these conditions are identical apart from the their framing and the inversion of the pecuniary incentives. This gives us suggestive evidence that framing effects are not playing a major role on determining search behaviour in the laboratory setting – confirmed in Table 5.2 below.

It is also noteworthy that the number of ‘never searchers’ is very similar in the *Search* and *SearchNoRej* conditions – consistent with the fact that for ‘never searchers’ these conditions are identical – since ‘never searchers’ never search, they do not experience rejection and therefore remain unaffected by the experimental variation which I have created in rejection across these conditions.

Figure 4: Distribution of search behaviour across treatment conditions

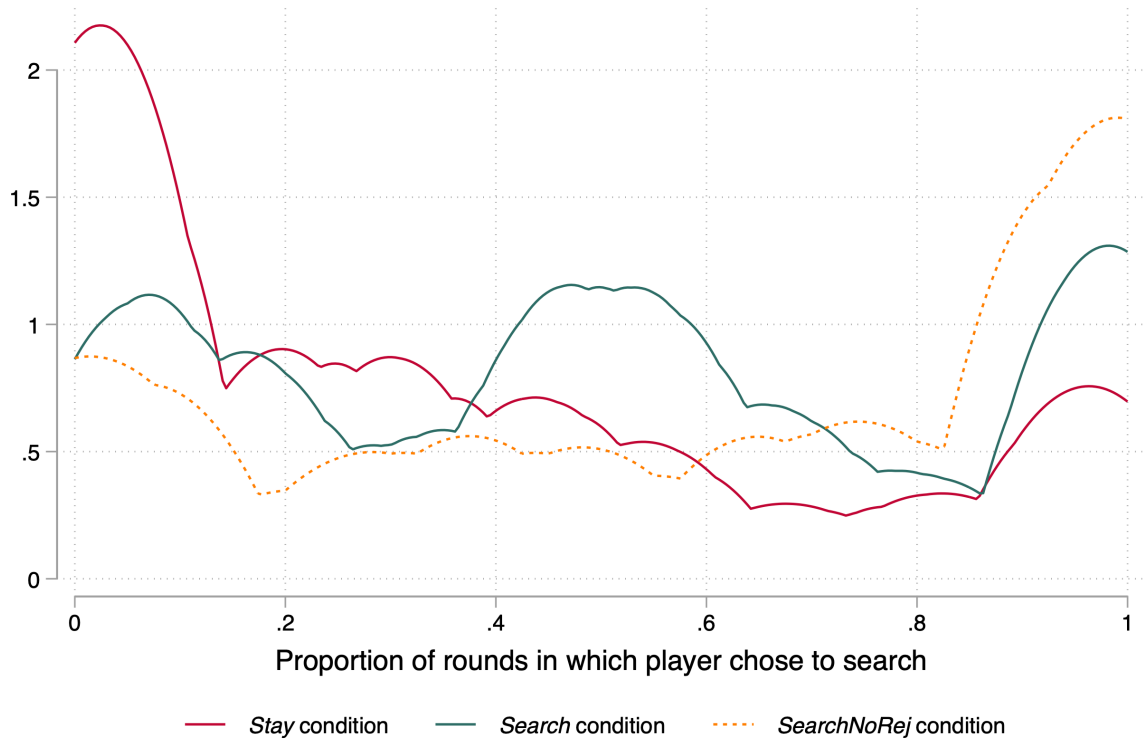


Table 5.1 reveals how this switching behaviour plays out across treatment conditions. The first row shows the search rates in Round 1, and the second shows the search rates over all rounds in the treatment condition (as shown in Figure 3). In the *Stay* condition, search rates begin low, at 36.5%, and remain stable throughout, at 32.8%. In the *Search* condition, search rates start high and then decline, while in the *SearchNoRej* condition, search rates start and remain high.

Across conditions, a substantial proportion of players do not switch strategies across rounds: Never switchers are most common in the *SearchNoRej* condition (59.5%), and lowest in the *Search* condition (39.4%), reflecting the fact that players are exposed to

rejection in *Search* but not in *SearchNoRej*.

In all conditions there is evidence that players experiment by ‘trying out’ other search conditions. Most often, these switches are subsequently reversed: That is, a player that starts out searching may choose in a subsequent round to ‘stay’, but then returns to ‘searching’. It is worth noting that in the *Search* condition, players are more likely than in either the *Stay* or the *SearchNoRej* condition to switch strategies terminally — i.e. to choose to ‘stay’ after initial choices to ‘search’ without then returning to ‘search’.

On average, players switch strategies for the first time between the third and the fourth round, and switch back to the original strategy between the fourth and the fifth round. The number of rounds in which players choose the alternate strategy before switching back to the original strategy is between 1.6 and 1.9 rounds.

These switching dynamics reveal that a substantial proportion of players choose and stick to a strategy, while others experiment with switching — regardless of the treatment condition.¹⁶ The average differences observed in search rates across treatment conditions are therefore likely driven by the increase in *terminal* switches within the *Search* condition.

Table 5.1: Search dynamics across treatment conditions

	Stay	Search	SearchNoRej
Searcher in R1 (%)	36.5	64.1	67.8
Proportion of rounds searched (%)	32.8	51.8	61.8
Switching behavior (%)			
Never switcher	47.3	39.4	59.5
Switcher (terminal)	15.8	24.1	5.0
Switcher (reversed)	37.0	36.5	35.5
Switch dynamics			
Round of first switch	3.2	3.6	3.3
Round of first return	4.7	4.6	4.9
Rounds between switch and return	1.8	1.6	1.9
Observations	203	170	121


¹⁶The presence of this reversible switching makes a Hazard model inappropriate for analysing search behaviour in this setup.

5.2 Performance on the real effort task

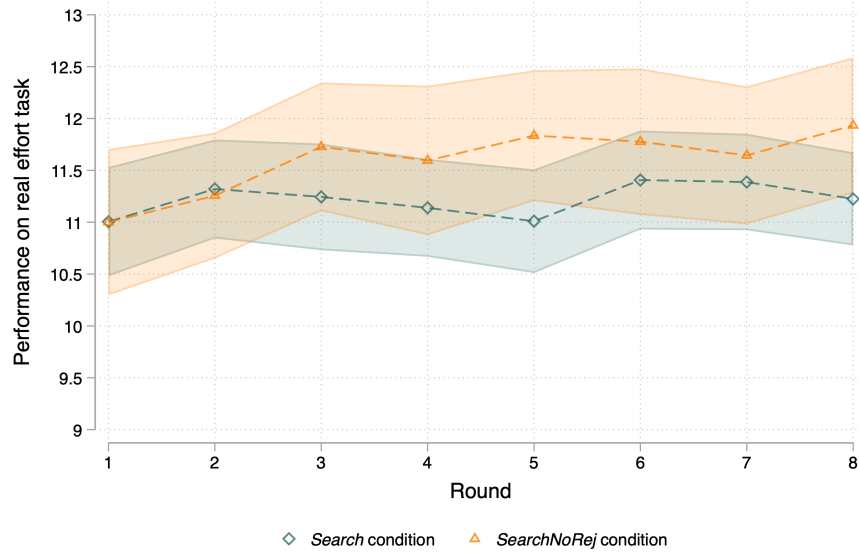
A secondary pre-specified outcome of the study was the effect of rejection on effort, as measured by performance on the real effort task. In the game, a discouraged worker might respond to rejection by adjusting behaviour on two margins: By choosing to search less, and/or by putting less effort into search. Because rejection in the game is a function of performance on the real effort task, we might be worried that if an initial rejection leads to a decline in effort, this may make subsequent rejections more likely. This would make it impossible to isolate the effect of rejection on search from the effects of declining effort. Figure 5 shows that effort remains stable over the rounds of the game, and does not differ across treatments. On average, participants completed between 11 and 12 tasks out of a possible 30.

Figure 5: Performance on the real effort task

(a) *Stay vs Search*



(b) *Search vs SearchNoRej*



Notes: Subjects could complete between zero and 30 slider tasks. The Vertical axis gives the number of tasks completed per round. Shaded areas represent the 95% confidence interval of point estimates in each round. $n = 179$ 109 for Figures a) and b) respectively.

5.3 Estimating the effect of rejection on job-search

In this subsection I present results of a series of fixed effects regressions to compare average search rates across the conditions discussed above. In all tables, the dependent variable is the binary ‘search’/‘stay’ outcome observed at each round r in each condition c and for each individual subject i . The independent variable of interest T is defined as a dummy equal to one if the individual was in the treatment condition of interest and equal to zero if the individual was in the comparison treatment condition.

The use of a within-subject experimental design requires at least two additional measures for causal identification: 1) Subject fixed effects A_i are included so as to control for time-invariant subject characteristics; 2) At a given point in time, a subject’s experience of rejection and job finding is determined by the choices she makes as well as the order in which she receives treatments. Thus, subjects in a given game may have had a different ‘history’ within the session and therefore will have experienced different amounts of rejection and/or job finding in prior games. For this reason, I also control for past outcomes in previous games within the session P_{ic-n} .

Specifications take the following form:

$$Search_{rci} = \beta_0 + \beta_1 T_i + A_i + P_{ic-n} + \epsilon_{rci} \quad (1)$$

Since the unit of randomisation is individual, but the unit of analysis is individual-time period, the error term ϵ_{rci} is clustered at the individual level (Abadie et al., 2017; Kim, 2022)

The order of treatments *Search* and *Stay* was randomised across experimental sessions, while *SearchNoRej* was implemented later in an add-on session. This means that ordering was not randomly varied between *Search* and *SearchNoRej*. For the sake of consistency, for this reason I exclude ordering fixed-effects in the main specifications. In a robustness test in Appendix Table E.1 I investigate whether including ordering fixed-effects has an impact on results, and find no evidence that it does.

In a first specification I run regressions on the full sample of subjects. In a second specification I limit the analysis only to those subjects who were never made an offer in any of the two conditions under consideration. This is because successful searchers

effectively select out of the sample since we no longer observe their choices once they receive an offer. In effect, this differential attrition may bias search rates downwards in the full sample relative to a specification which excludes job finders. However, since the same dynamic affects both the treatment condition of interest and the comparison condition, I do not expect this to be a consequential distortion.

Table 5.2: Treatment effects of rejection exposure on search rates

	% of players searching in given round			
	(1)	(2)	(3)	(4)
Exposure to rejection (<i>Search</i>)	-0.148*** (0.048)	-0.182*** (0.051)	-0.155*** (0.051)	-0.146*** (0.052)
Constant	0.685*** (0.022)	0.708*** (0.022)	0.661*** (0.030)	0.657*** (0.032)
Mean search rate, no exposure to rejection	0.69	0.69	0.69	0.69
% difference in search rates	-21.42	-26.36	-22.57	-21.27
Number of observations	3143	2944	2537	2408
Number of players	203	196	180	173
Indiv. Fixed Effects	Yes	Yes	Yes	Yes
Past outcomes in session	Yes	Yes	Yes	Yes
Framing effects removed (<i>SearchNoRej</i>)	No	No	Yes	Yes
Excludes job-finders	No	Yes	No	Yes

Notes: * $p < .1$; ** $p < .05$; *** $p < .01$. Standard errors, clustered at the subject level, are in parentheses. The outcome variable in both specifications is a binary variable equal to 1 if a subject decides to ‘search’ in a given round, and equal to 0 if a subject decides to ‘stay’. The independent variable of interest, *Treatment condition with exposure to rejection* is defined as a dummy equal to one if the individual was in the treatment condition of interest (*Search*) and equal to zero if the individual was in the comparison treatment condition (*Stay (inverted)* in the first two columns, and *SearchNoRej* in the last two columns). Coefficients can be interpreted as the treatment effect on search rates of being in the *Search* condition relative to the comparison condition (in which subjects are not exposed to rejection in the game). The ‘Full sample’ columns report results for the full sample, while the ‘Excluding job finders’ columns report results only for the subsample of subjects who did not search and receive a job offer in either condition under consideration.

Regression results confirm the narrative which emerges from the preceding analysis: Compared to what we would expect search rates to be if subjects were indifferent to receiving information on rejection, in Table 5.2 we observe search rates to be significantly lower in reality when subjects are exposed to information on rejection. I interpret this difference as being driven by rejection information avoidance behaviour: Search rates in

the *Search* condition are between 15 and 18 percentage points (or 21-26 percent) lower than what we would expect in the absence of rejection aversion.

Note that, as in Figure 3 above, in the second two columns of Table 5.2 I am able to isolate the rejection effect from possible framing effects by comparing search rates in the *Search* condition to those in the *SearchNoRej* condition. Having netted out possible framing effects, the effect of rejection on search rates persists and remains substantively unchanged. In Appendix F I provide a direct test for framing effects, and find that framing effects are small and statistically insignificant, with null effects fairly tightly centered on zero.

5.4 Correlates of search behaviour and discouragement

Appendix H investigates the demographic and behavioural correlates of the search behaviour I observe in the results presented above. Overall, I observe that older individuals are less likely to switch away from Search following a rejection, while black African workers are more likely than the coloured subsample to switch away from Search following a rejection. Interestingly, gender does not appear to be correlated with search behaviour in the laboratory game.¹⁷ Those individuals who report struggling with their motivation to search for jobs in the real labour market also appear to be more easily discouraged in the game. There is some weak evidence of a correlation between a measure of self-assurance ('attributing success to one's own merit') and persistence in Search. Ranking above average on a self-efficacy scale (compared to below average) yields insignificant results. Importantly, I also find that individuals who are very optimistic about their chances of being retained in their (real) jobs, are also more likely to switch away from Search following a rejection.

It is difficult to determine whether the null result on self-efficacy is a true null, whether this is because such variables are difficult to reliably measure, or whether a deeper issue may be at play: Specifically, as suggested by work on fragile self-esteem (Kőszegi et al., 2021), we would expect individuals to avoid information on rejection if they have high but

¹⁷Others have suggested that gender may be an important correlate of discouragement, with women more likely to be discouraged by rejection than men (Bapna et al., 2021).

fragile self-esteem.¹⁸ While it may be possible to measure self-esteem, it is more difficult to measure the fragility of these beliefs. This gets in the way of a straightforward mapping of self-esteem beliefs onto rejection avoidant behaviour, and may help explain such null results.

6 A simple decision-making framework accounting for preferences for information on rejection

In the laboratory experiment I describe above, I find that: *a)* subjects appear to be initially naive with regard to their aversion to receiving information on rejection, *b)* subjects respond to early rejections by switching to an activity which has lower expected pecuniary returns but which also has lower exposure to high-frequency information on rejection, and *c)* that the drop-off in average search rates is steepest after early rejections and subsequently flattens out. In this section I propose a basic theoretical framework to help illuminate the mechanisms underlying these findings.

Let y_{itj} denote whether a positive payoff is received for individual i in each period t in activity j . The probability of obtaining a payoff in an activity is given as $p_{it} = e_{it} + \epsilon_{it}$, where e_{it} denotes effort by individual i , $0 \leq e_i \leq 1$, in time t and $\epsilon_{it} \sim N(0, \sigma^2) \forall i, t$. In addition, the payoff, the value of which is denoted by wage w_{itj} , is positive iff $p_{it} > \tau$ for some exogenous threshold τ .

¹⁸In Kőszegi et al. (2021)’s model, individuals have multiple ‘personal self-esteem equilibria’, and are either consciously or subconsciously aware that shocks to self-esteem may shift them from one equilibrium to another. Of particular relevance are *fragile* high self-esteem equilibria – characterised by a large distance between positively biased beliefs about ability and true (low) ability. Kőszegi et al. (2021) show that these individuals face especially strong incentives to distort their choices so as to avoid a downward revision of self-esteem. Applying this model to job search, Kőszegi et al. (2021) propose that their model helps resolve an apparent paradox: One the one hand, the unemployed are systematically over-optimistic about finding jobs (Spinnewijn, 2015), while on the other hand, they appear to search surprisingly little (A. B. Krueger and Mueller, 2010; A. Krueger and Mueller, 2012). According to the ‘fragile self-esteem’ model, this paradox can be explained by the fact that individuals derive utility from occupying a high self-esteem equilibrium, but being aware that this equilibrium is ‘fragile’, avoid exposure to shocks – for instance, rejections in job search – which could move them on to a lower equilibrium. Kőszegi et al. (2021) imply that individuals may be willing to pay – in the form of foregone income and employment – to avoid these self-esteem shocks.

Thus

$$y_{it} = \begin{cases} 1 & \text{if } p_{it} > \tau \\ 0 & \text{otherwise} \end{cases}$$

Hence the wage $w_{itj} = w_j \cdot y_{it}$, $j \in \{J, S\}$

Rearranging, we can see that $p_{it} > \tau$ can be expressed as $(e_{it} - \tau) + \epsilon_{it} > 0$. Expressed in this way, and given the distributional assumptions made above, we can express $E(y_{it})$ as a probit model, where

$$Pr(y = 1|p_{it}) \equiv \Phi(e_i - \tau). \quad (2)$$

I assume that individual i draws utility from wages w_{itj} and disutility from experiencing rejection ρ_{it} , both of which are informed by activities they undertake in each period. Individuals choose between two mutually exclusive activities $j \in \{J, S\}$ (as in the experiment) in each time period t : J , denoting effort in the current job and S , denoting effort in search for another job. The additively-separable intertemporal utility function for individual i in time t in activity j is given by:

$$U_j(w_{itj}, \rho_{it}) = w_{itj} - \lambda_i(\alpha_{i0} \cdot \rho_{it})^2 \quad (3)$$

where

$$\rho_{it} = \sum_{t=1}^T (1 - y_{it}). \quad (4)$$

w_{itj} denotes individual i 's payoff in activity j . The weight parameter λ_i captures the extent to which individual i has a preference for not receiving information on rejection. If $\lambda_i = 0$, we return to the usual framework where preferences are determined over monetary returns w_{tj} alone. ρ_{it} is the cumulative number of rejections an individual has received up to time t . ρ_{it} is multiplied by α_{i0} , representing an individual's initial beliefs about ability (i.e. prior to receiving any rejection), so that the disutility of rejection increases

in α_{i0} . This captures the notion that individuals with elevated beliefs about ability have ‘more to lose’ with experiences of rejection (Kőszegi et al., 2021). Note too that I use a quadratic form for the ‘disutility of rejection’ term, so as to capture increasing marginal costs of rejection.¹⁹

In line with the set up of the experiment, I assume that in activity S (the Search option), the individual is exposed to rejection after completing a task. In contrast, in activity J (the stay option), the individual does not receive rejection signals. I also assume that activity S yields higher returns than activity J in pecuniary terms, so that $w_S > w_J \forall t$.

Letting Φ denote $\Phi(e_i - \tau)$ (see equation 2 above), the expected utility in each activity can then be written as:

$$\begin{aligned}
S : EU_S(w_{itS}, \rho_{it}) &= E(w_{itS}) - E(\lambda_i(\alpha_{i0} \cdot \rho_{it})^2) \\
&= w_S \cdot E(y_{it}) - \lambda_i \alpha_{i0}^2 E(\rho_{it}^2) \\
&= w_S \Phi - \lambda_i \alpha_{i0}^2 E(\rho_{it-1} + (1 - y_{it}))^2 \\
&= w_S \Phi - \lambda_i \alpha_{i0}^2 E((\rho_{it-1} + 1)^2 + 2(\rho_{it-1} + 1)y_{it} + y_{it}^2) \\
&= w_S \Phi - \lambda_i \alpha_{i0}^2 (\rho_{it-1} + 1)^2 + 2(\rho_{it-1} + 1)\Phi + \Phi
\end{aligned}$$

$$\begin{aligned}
J : EU_J(w_{itJ}, \rho_{it}) &= E(w_{itJ}) - E(\lambda_i(\alpha_{i0} \cdot \rho_{it})^2) \\
&= w_J \cdot E(y_{it}) - \lambda_i (\alpha_{i0} \cdot \rho_{it-1})^2 \text{ since in J } \rho_{it} = \rho_{it-1} \\
&= w_{tJ} \Phi - \lambda_i \alpha_{i0}^2 (\rho_{it-1}^2)
\end{aligned}$$

Simplifying, the choice of activity in each period can then be summarised by:

$$\begin{aligned}
\text{Activity} &= S \text{ if } (w_S - w_J)\Phi > \lambda_i \alpha_{i0}^2 (2\rho_{it-1} + 1 + 2(\rho_{it-1} + 1)\Phi + \Phi) \\
&J \text{ otherwise}
\end{aligned} \tag{5}$$

Thus, individuals are more likely to search if λ_i is low, ρ_{it} is low, and/or $w_{tS} - w_{tJ}$ is high.

¹⁹This quadratic form captures the notion that the disutility if a later signal is larger than that of an earlier signal, as individuals get ‘worn down’ by rejection and become increasingly likely to avoid information following ‘bad’ news (Karlsson et al., 2009).

Note also that in the final period subjects will receive information on the outcome of their choice regardless of which activity they choose. For this reason, the ‘disutility of rejection’ term in EU_S and EU_J are identical in the final period, and only differences in earnings should affect choices. Since $w_{iS} > w_{iJ} \forall t$, individuals should revert to search in the final period.


Figure 6 graphically illustrates some of the characteristics of the model by plotting the search behaviour of two paradigmatic individuals over eight rounds of the game. In the representation in Figure 6, the model is parameterised so that in Round 1 both individuals search – that is, the pecuniary benefits of ‘searching’ dominate the non-pecuniary benefits of ‘staying’. These two paradigmatic individuals are identical across all parameters apart from those that are specified to be different across the two individuals – λ_i in Figure 6a and α_{i0} in Figure 6b.

Individuals will switch from choosing Option S (‘Searching’) to Option J (‘Staying’) at the point at which the non-pecuniary benefits of ‘staying’ (the right hand side of expression 5) switch from being dominated by the pecuniary benefits to searching to dominating the pecuniary benefits of searching.


We can see in Figure 6a that individuals with low λ_i values – i.e. a low weight placed on rejection in the utility function – will persist in search longer than individuals with higher λ_i values. This makes intuitive sense – it will take more rejection for the pecuniary benefits of searching to be dominated by self-esteem concerns among individuals who place relatively little weight on rejection. Figure 6b illustrates that individuals with higher values of initial self-esteem (α_{i0}) will drop out of search sooner than individuals with initially lower levels of self-esteem. This is because the distance between the rejection signal and initial beliefs is larger for high self-esteem individuals than for low self-esteem individuals.

Figure 6: Comparing job search behaviour between individuals with different parameter values

a) High versus low λ_i values (preferences to avoid rejection)



b) High versus low α_{i0} values (initial level of self-esteem)



Notes: The Figure illustrates the search behaviour of two paradigmatic individuals over eight rounds of the game. Players search if $(w_S - w_J)\Phi > \lambda_i \alpha_{i0}^2 (2\rho_{it-1} + 1 + 2(\rho_{it-1} + 1)\Phi + \Phi)$. The model is parameterised so that in Round 1 both individuals search – that is, the pecuniary benefits of ‘searching’ dominate the non-pecuniary benefits of ‘staying’. Apart from differing in the specified parameter values in each respective figure, the two paradigmatic individuals are identical across all other parameters. Individuals will switch from choosing Option S (‘Searching’) to Option J (‘Staying’) at the point at which the non-pecuniary benefits of ‘staying’ (the right hand side of the above expression) dominate the pecuniary benefits to ‘searching’. Round 8 is assumed to be the final round. As described above, I show that all players ought to switch to Search in the final round.

It is worth noting a few limitations of the model: First, the model does not explicitly account for the *fragility* of self-esteem beliefs. Following Kőszegi et al. (2021), we might expect that the fragility of self-esteem beliefs will be an important determinant of an individual’s propensity to actively avoid information that is not consonant with prior beliefs.

Second, the model does not account for the role of salience: Following Bordalo et al. (2022), we can contrast ‘stimulus salience’ as a *bottom-up* mechanism driving behaviour to ‘goal relevance’ as a *top-down* mechanism driving behaviour. In the case of the game played in the laboratory experiment, the ‘goal’ is to maximise earnings in the game, while (repeated) rejection may function as a salience stimulus in which ‘bottom-up attention interferes with goals and causes suboptimal decisions’ (Bordalo et al. (2022), p. 525). Salience simply changes the *weight* placed on a particular stimulus. In this case, the disutility of experiencing rejection may be weighted more heavily when it is more salient in high-rejection conditions. Importantly, this does not change the *direction* in which the stimulus pulls, but merely the weight placed on the stimulus. The role of salience could be formally accounted for in the basic model described above by multiplying the ρ_{it} rejection disutility term by a salience-weight parameter, and allowing this weight to increase with past exposure to rejection.

For the sake of simplicity and tractability, neither the role of salience nor the role of the fragility of prior beliefs is accounted for in this basic model.

7 Discussion

7.1 Discussion of main results

Result 1: Prior to receiving any rejections in a given game, the average subject chooses the option that maximises the expected payoff rather than that which minimises exposure to rejection information.

In the first round of the game, the average subject appears to respond to monetary incentives as if they have no preferences for information – that is, observed behaviour (see Figure 3a) approximates the paradigmatic cases presented in Figure 6. In other words, in the first round of a game, most subjects choose to ‘search’ when there is a monetary

incentive to do so, and most choose to ‘stay’ when there is a monetary incentive to do so.

While in all games one option always strictly dominates the other in terms of expected monetary returns, the absolute difference in expected returns is very small once we account for the low probability of receiving a payoff. Therefore, if we account for some randomness in individuals’ decision making, it is not clear ex ante that individuals would respond to the (small) pecuniary incentives in the way that we observe they do, even putting aside preferences for information.

It therefore seems likely that individuals are engaging in heuristic reasoning to guide their decisions prior to exposure to rejection, where the simple rule of thumb is ‘choose the option with the highest expected pay-off’. Two features of the game make the use of such heuristics more likely (Tversky and Kahneman, 1974): First, subjects were encouraged to make decisions quickly: The other subjects in the session would only be able to proceed at the pace of the slowest subject, which created pressure on all subjects to make quick decisions. Second, before beginning each game subjects were required to correctly identify in which of the two tasks ‘you can earn more’. While expected differences in earnings between the two tasks is very small in absolute terms, by prompting subjects to simply identify which of the two is more remunerative, I potentially increase the propensity for subjects to default to the higher-payoff activity. As long as individuals are engaging in this heuristic reasoning, the role of random decision-making errors may be fairly weak, even if the absolute difference in expected returns between options is small.

Result 2: Exposure to rejection leads to a rapid and dramatic reduction in search rates: The average subject shifts from choosing the option that maximises the expected payoff to the option that minimises exposure to rejection information (when these are not the same).

Subjects respond to receiving rejection information by switching to an activity that has lower expected pecuniary returns but that also has lower exposure to high-frequency information on rejection. Since the disutility of rejection term in Equation 4 is not fixed but increases with exposure to new information on rejection, we could expect to see individuals become more information avoidant after exposure to information which is discordant with their initial beliefs.

If this is the case, we may expect to see deviations from heuristically-based reasoning as information avoidance increases in importance within the individual’s utility function

relative to consumption utility. For some individuals the disutility of rejection will come to dominate the heuristic which dictates ‘choose the option with the highest pecuniary return’. This may precipitate a switch to the activity which has lower expected pecuniary returns but which also has lower exposure to high-frequency information on rejection. As shown in Figure 6, we would expect that individuals who place a higher weight on avoiding rejection and individuals with higher initial self-esteem beliefs will switch earliest.

Result 1 above implies that the anticipated costs of rejection information may initially be small as agents are confident about outcomes and/or may not know how sensitive their beliefs will be to new information. However, as rejection is repeated – representing compound shocks to self-esteem – the anticipated costs of risking further rejection may increase, leading to an increase in active information avoidance.

Result 3: The drop-off in search rates is steepest after the first rejections

The drop-off in average search rates is steepest after early rejections and subsequently flattens out. In the framework above I assume that individuals are heterogeneous in terms of their preferences for avoiding rejection and in terms of their initial beliefs about their ability. As graphically illustrated in Figure 6a, this implies that a large negative shock to self-esteem will not enter strongly into an individual’s utility function if the weight placed on self-esteem is small or if initial beliefs about ability are modest. The upshot of this is that some individuals will be slower to stop searching than others. The patterns observed in Figures 3c and 4 suggest that the drop-off in search rates over the rounds of the game is driven by selection out of search on the part of highly rejection-sensitive individuals.

7.2 Alternative explanations for observed behaviour

The identification of an effect of information avoidance on search behaviour hinges on being able to rule out alternative explanations. In this subsection I discuss several of these alternative potential drivers of search behaviour and the concerns they might pose to clean identification.

7.2.1 Redirecting effort in response to (perceived) signals about comparative advantage

In the real world, rejection may carry instrumentally valuable information which may help job-seekers direct their search more profitably or refine their approach to signalling employability. For instance, (repeated) rejection in one type of activity may signal to a job-seeker that her comparative advantage is elsewhere and so may motivate her to redirect her search or otherwise change strategies – which may include focusing on her current job.

This poses a potential problem to this study: If subjects in the game believe that they are learning instrumentally valuable information from rejections, when they choose to search and are rejected they may (mistakenly) interpret rejection as a sign that they have a higher chance of receiving an offer in the ‘stay’ task.. Therefore, if subjects believe that they are learning about the returns to search through rejections, this may explain some of the switching that I observe from the ‘search’ to the ‘stay’ tasks. In reality, the likelihood of receiving an offer in each task is identical, conditional on effort.

I take several measures to minimise the possibility of a subject making the mistake of believing that the likelihood of receiving a payoff in the ‘search’ task is lower than in the ‘stay’ task. I am fully transparent about the way in which outcomes are determined, emphasising that the underlying outcome generating process is held constant across both tasks. I provide high frequency information on subjects’ performance, the likelihood of success, and the relationship between performance and the likelihood of success in a way that is *identical* regardless of whether a subject chooses to ‘search’ or ‘stay’. This is a major focus of the session introduction and the regular comprehension checks (see Appendix C). In addition, the real effort task is the same for both the ‘search’ and the ‘stay’ activities, which eliminates the possibility of subjects believing that they are better at one task than the other. In both activities, immediately after the real effort task I provide information on performance (i.e. 12/30 tasks completed) and the corresponding likelihood of receiving an offer (4 percent chance of receiving an offer). This helps emphasise that the relationship between effort and outcomes is *identical* whether one chooses to ‘search’ or to ‘stay’. The only thing that differs between conditions is the information that is provided on the realisation of the outcome – which in itself carries no information on the mechanisms determining outcomes in the game beyond that which is already provided.

Of course, it remains possible that subjects do not fully understand that the relationship between effort and outcomes is identical in the two tasks, or that they do not believe that they are. Beyond the measures that I have taken, I cannot definitively rule out this possibility.

7.2.2 Relevance of self-esteem beliefs in a quasi-lottery

Outcomes in the game are always primarily determined by luck. On a round-by-round basis, the probability of receiving an offer varies from 0-10 percent.²⁰ A bad outcome is therefore only very marginally related to performance on the real effort task. One may wonder, therefore, whether failure would be perceived as a shock to self-esteem, or merely as a self-esteem-irrelevant disappointment in the same way that losing at a lottery might be disappointing but not hurtful. If this is the case, the information avoidance effect we observe simply collapses into disappointment aversion, rather than the more specific information-on-rejection avoidance mechanism I have proposed.

Rejection falls under the broader category of disappointment. To be rejected is to be disappointed, while to be disappointed is not necessarily to be rejected. Thus, the presence of disappointment aversion is not a threat to the causal story I tell as long as a part of this disappointment is self-esteem-relevant. While it is hard to disentangle rejection from disappointment, being unable to do so is inconsequential from a policy perspective, since the implications are the same.

An observation (albeit anecdotal) from the experimental sessions provides evidence that that subjects systematically over-weighted the role of performance relative to luck in determining outcomes – suggesting a role for self-esteem-relevant beliefs: At the beginning of the experimental session I personally went through a detailed set of instructions and explanations, which included a clear and precise explanation of how outcomes were determined partly by a subject’s performance on the real effort task, and partly by luck. Before starting each game, subjects needed to complete a comprehension test. One of the questions was ‘What determines whether you get offered a job or not?’, and subjects were given the response options: ‘How many of the slider tasks you complete in a given round’, ‘Luck’, and ‘Both answers’ (See Appendix C). Subjects were only able to proceed

²⁰In practice, the probability of receiving an offer varied from 0-8 percent since the best score on the real effort task (24/30) corresponds to an 8 percent chance of receiving an offer.

once they had answered all questions correctly. While supervising this comprehension test, I observed that the overwhelming majority of subjects got stuck on this question – answering that ‘How many of the slider tasks you complete in a given round’ determines game outcomes, rather than the correct answer that both luck and performance determines outcomes.²¹ This suggests that, despite explicit explanations to the contrary, subjects systematically placed too much emphasis on the role of performance, and too little emphasis on the role of luck. While the comprehension test may have partially corrected this tendency to over-weight performance, it is likely that this phenomenon helps explain the strength of the behavioural response to rejection in the game.

7.2.3 Errors in probabilistic reasoning

I find no evidence that errors in *reasoning* are driving the discouragement results. Despite the comprehension checks and explicitly stating that probabilities of receiving an offer are identical whether a subject chooses to ‘search’ or ‘stay’, one might still worry that subjects erroneously reason that a sequence of rejections may exhibit positive recency (i.e. that a rejection is more likely to follow a sequence of rejections) or may exhibit negative recency (i.e. that an offer is more likely to follow a sequence of rejections). Positive recency would be especially concerning, since a switch away from ‘search’ following a string of rejections might be motivated by an erroneous belief that a rejection is more likely to follow a rejection. If this is the case, this positive recency might be confused for information avoidance.

In the post-session survey I measured subjects’ propensity to commit hot hand fallacies (positive recency) or gambler’s fallacies (negative recency).²² One would be concerned if subjects who display a propensity to commit hot hand fallacies / display positive recency are less likely to ‘search’, since this may be a potential confounder which would undermine my ability to make causal claims about information avoidance. I see no evidence of this:

²¹Unfortunately, I did not anticipate responses to these comprehension checks to be of analytical interest, and so did not record responses and so I cannot give a quantitative sense of the relative frequency of these errors.

²²I used two questions motivated by Ayton and Fischer (2004)’s finding that ‘where subjects construe a sequence of outcomes as reflecting human performance [they demonstrate] positive recency [while] where subjects expect that outcomes are due to an inanimate mechanism [they demonstrate] negative recency.’ The two questions I used were: 1) ‘Imagine you flip a fair coin, and get heads five times. How likely is it that it will fall on heads when you flip it a sixth time?’, and, 2) ‘Imagine you ask a friend to guess a coin flip, and flip it five times. The results are: H, H, T, H, T. Your friend guesses right every time! How likely is it that she will get it right the sixth time?’.

If anything, those subjects who have a propensity to commit hot hand fallacies appear *more* likely to search.²³ This suggests that the observed rejection effect is robust to the potential presence of gambler fallacy / hot-hand fallacy-type reasoning.

8 Conclusion

I use a behavioural game implemented with young workers in South Africa to study the effect that rejection signalling has on job search. In a multi-round game I ask subjects to make a choice between two activities – to ‘search’ or to ‘stay’. At the end of the game, subjects will either have found a ‘new job’ by ‘searching’, have been offered a ‘contract extension’ by ‘staying’, or neither. The two activities differ in the (known) probability that the subject will receive a rejection message after each round but not in the (known) overall likelihood of receiving a job offer (whether current or new). By experimentally manipulating the expected returns (in terms of compensation incentives) and the frequency of rejection signals across treatment conditions, I am able to isolate the effect of rejection signalling on job search, holding potential confounders, including framing effects, constant.

Two main results stand out. First, subjects display strong rejection aversion and are willing to trade-off expected earnings in order to reduce their exposure to rejection. This results in a reduction in search relative to search rates of a counterfactual rejection-indifferent agent. These results are driven by the rejection messaging, rather than by framing effects.

Second, the intensity of this effect increases with exposure to rejection. The design of the game allows me to interpret rejection-avoidant behaviour as a form of active information avoidance. The trends I observe in rejection avoidant behaviour show that this information avoidance is not a static phenomenon, but changes as people are exposed to rejection. To begin with, subjects respond to pecuniary incentives as we would expect them to if they were indifferent to information on rejection. However, I find that past exposure to rejection makes people more willing to pay to avoid further exposure to

²³The average search rate for those who have a propensity to commit hot hand fallacies is 53 percent, compared to 51 percent for those who do not display a propensity to commit hot hand fallacies. The proportion of those who stop searching following a rejection is also similar – at 35 percent and 34 percent respectively. These differences are not statistically significant.

rejection.

These findings have implications for mechanisms to increase job search and on-the-job search: Cases in which applications are due approximately at the same time and where outcomes are communicated after all applications are submitted may increase the total number of applications submitted, compared to cases in which search is more ‘sequential’ in the sense that both search and outcomes are continuous. Examples of the former include, for instance, graduate school applications and some academic job markets, while examples of the latter are more typical in low-wage white and blue collar work.

In this paper I frame the discussion in terms of job search. However, this framework could be extended to any domain in which individuals face choices between high-rejection and low-rejection activities and where information on rejection carries a psychological cost that might affect behaviour. Of particular interest are classes of activities in which choices are made between high-rejection frequency, high-reward and low-rejection frequency, low-reward activities. Dating and university applications may fall into this category. For instance, the potential returns to being accepted into a top-ranked university may be very high, but applications to these universities carry with them a high potential for rejection. Rejection aversion may create disincentives to applying to top schools in favour of less competitive schools, with real welfare consequences for the individual. In such cases a similar framework as the one used in this paper might be usefully applied.

Online dating applications have risen in popularity in part because of innovations in how they have dealt with the hurt of rejection. In the most popular applications, a user requests to match with another user through the centralised platform. The cost to requesting a match is negligible, merely the flick of a thumb. The platform will inform the user only if a match is successful. Users are not informed if a match is unsuccessful. This silence is deliberately ambiguous, since rejection is only one possible explanation—the request may not have been seen (yet) by the other user, or the other user may be inactive.

In the job market, the emergence of online platforms and centralised and/or coordinated job markets may already go some way towards reducing the psychological burden of rejection, at least during the job search phase. In future, job matching platforms may be successful in increasing engagement by adopting strategies inspired by the online dating sector: Not only by lowering the costs to making applications, but also by controlling the information that is (not) shared with applicants on outcomes. A barrier to this approach in the past has been the cost burden it would place on employers to screen applicants.

However, thanks to recent technological advances, artificial intelligence is increasingly being used to cheaply and effectively screen job applicants. Using such innovations to centralise job search and to remove psychological frictions may help improve labour market outcomes for job seekers, and improve the efficiency with which talent is allocated to jobs in general.

References

- [1] Abadie, Alberto, Athey, Susan, Imbens, Guido W, and Wooldridge, Jeffrey. *When Should You Adjust Standard Errors for Clustering?* Working Paper 24003. National Bureau of Economic Research, Nov. 2017. DOI: 10.3386/w24003. URL: <http://www.nber.org/papers/w24003>.
- [2] Abel, Martin, Burger, Ronelle, Carranza, Eliana, and Piraino, Patrizio. “Bridging the intention-behavior gap? The effect of plan-making prompts on job search and employment”. In: *American Economic Journal: Applied Economics* 11.4 (2019), pp. 284–301.
- [3] Andolfatto, David, Mongrain, Steeve, and Myers, Gordon. “Rational truth-avoidance and self-esteem”. In: *Canadian Journal of Economics/Revue canadienne d’économique* 42.1 (2009), pp. 141–154.
- [4] Andreoni, James, Rao, Justin M, and Trachtman, Hannah. “Avoiding the ask: A field experiment on altruism, empathy, and charitable giving”. In: *Journal of political Economy* 125.3 (2017), pp. 625–653.
- [5] Andries, Marianne and Haddad, Valentin. “Information aversion”. In: *Journal of Political Economy* 128.5 (2020), pp. 1901–1939.
- [6] Ayton, Peter and Fischer, Ilan. “The hot hand fallacy and the gambler’s fallacy: Two faces of subjective randomness?” In: *Memory & cognition* 32.8 (2004), pp. 1369–1378.
- [7] Banerjee, Abhijit, Galiani, Sebastian, Levinsohn, Jim, McLaren, Zoe, and Woolard, Ingrid. “Why has unemployment risen in the new South Africa? 1”. In: *Economics of Transition* 16.4 (2008), pp. 715–740.
- [8] Bapna, Sofia, Benson, Alan, and Funk, Russell. “Rejection Communication and Women’s Job-Search Persistence”. In: *Available at SSRN 3953695* (2021).

- [9] Berg, Gerard J van den and Uhlendorff, Arne. “14 Economic Job Search and Decision-Making Models”. In: *The Oxford Handbook of Job Loss and Job Search* (2018), p. 243.
- [10] Bjørnstad, Roger. “Learned helplessness, discouraged workers, and multiple unemployment equilibria”. In: *The Journal of Socio-Economics* 35.3 (2006), pp. 458–475.
- [11] Bordalo, Pedro, Gennaioli, Nicola, and Shleifer, Andrei. “Salience”. In: *Annual Review of Economics* 14 (2022), pp. 521–544.
- [12] Burns, Justine, Godlonton, Susan, and Keswell, Malcolm. “Social networks, employment and worker discouragement: Evidence from South Africa”. In: *Labour Economics* 17.2 (2010), pp. 336–344.
- [13] Caria, Stefano, Orkin, Kate, Andrew, Alison, Garlick, Roc, Heath, Rachel, and Singh, Niharika. “Barriers to Search and Hiring in Urban Labour Markets”. In: Working Paper Series 10(1) (Feb. 2024).
- [14] Charness, Gary, Gneezy, Uri, and Henderson, Austin. “Experimental methods: Measuring effort in economics experiments”. In: *Journal of Economic Behavior & Organization* 149 (2018), pp. 74–87.
- [15] Charness, Gary and Kuhn, Peter. “Lab labor: What can labor economists learn from the lab?” In: *Handbook of labor economics* 4 (2011), pp. 229–330.
- [16] Chen, Daniel, Schonger, Martin, and Wickens, Chris. “oTree—An open-source platform for laboratory, online, and field experiments”. In: *Journal of Behavioral and Experimental Finance* 9 (2016), pp. 88–97.
- [17] Chen, Gilad, Gully, Stanley M, and Eden, Dov. “Validation of a new general self-efficacy scale”. In: *Organizational research methods* 4.1 (2001), pp. 62–83.
- [18] Dana, Jason, Weber, Roberto A, and Kuang, Jason Xi. “Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness”. In: *Economic Theory* 33.1 (2007), pp. 67–80.
- [19] Donovan, Kevin, Lu, Will Jianyu, Schoellman, Todd, et al. *Labor market dynamics and development*. Tech. rep. JSTOR, 2020.
- [20] Field, Erica, Garlick, Robert, Subramanian, Nivedhitha, and Vyborny, Kate. “Why Don’t Jobseekers Search More? Barriers and Returns to Search on a Job Matching Platform”. Unpublished manuscript. Feb. 2023.

- [21] Ganguly, Ananda and Tasoff, Joshua. “Fantasy and dread: The demand for information and the consumption utility of the future”. In: *Management Science* 63.12 (2017), pp. 4037–4060.
- [22] Garlick, Robert, Hensel, Lukas, Kiss, Andrea, and Orkin, Kate. “Jobseekers’ Beliefs about Comparative Advantage and (Mis)Directed Search”. Unpublished manuscript. Mar. 2023.
- [23] Gill, David and Prowse, Victoria. “A structural analysis of disappointment aversion in a real effort competition”. In: *American Economic Review* 102.1 (2012), pp. 469–503.
- [24] Gill, David and Prowse, Victoria. “Measuring costly effort using the slider task”. In: *Journal of Behavioral and Experimental Finance* 21 (2019), pp. 1–9.
- [25] Golman, Russell, Hagmann, David, and Loewenstein, George. “Information avoidance”. In: *Journal of economic literature* 55.1 (2017), pp. 96–135.
- [26] Huffman, David, Raymond, Collin, and Shvets, Julia. “Persistent overconfidence and biased memory: Evidence from managers”. In: *American Economic Review* 112.10 (2022), pp. 3141–75.
- [27] Ingle, Kim and Mlatsheni, Cecil. *The extent of churn in the South African youth labour market: Evidence from NIDS 2008-2015*. Working Paper 201. Southern Africa Labour and Development Research Unit, 2017.
- [28] Jäger, Simon, Roth, Christopher, Roussille, Nina, and Schoefer, Benjamin. *Worker Beliefs About Outside Options*. Working Paper 29623. National Bureau of Economic Research, Dec. 2021.
- [29] Karlsson, Niklas, Loewenstein, George, and Seppi, Duane. “The ostrich effect: Selective attention to information”. In: *Journal of Risk and uncertainty* 38 (2009), pp. 95–115.
- [30] Kerr, Andrew. “Job flows, worker flows and churning in South Africa”. In: *South African Journal of Economics* 86 (2018), pp. 141–166.
- [31] Kim, Duk Gyoo. “Clustering Standard Errors at the Session Level”. In: *SSRN* (2022). URL: <https://dx.doi.org/10.2139/ssrn.4139997>.

- [32] Kőszegi, Botond, Loewenstein, George, and Murooka, Takeshi. “Fragile Self-Esteem”. In: *The Review of Economic Studies* 89.4 (Sept. 2021), pp. 2026–2060. ISSN: 0034-6527. DOI: 10.1093/restud/rdab060. eprint: <https://academic.oup.com/restud/article-pdf/89/4/2026/44998593/rdab060.pdf>. URL: <https://doi.org/10.1093/restud/rdab060>.
- [33] Krueger, Alan and Mueller, Andreas. “The lot of the unemployed: a time use perspective”. In: *Journal of the European Economic Association* 10.4 (2012), pp. 765–794.
- [34] Krueger, Alan B and Mueller, Andreas. “Job search and unemployment insurance: New evidence from time use data”. In: *Journal of Public Economics* 94.3-4 (2010), pp. 298–307.
- [35] Momsen, Katharina and Ohndorf, Markus. “When do people exploit moral wiggle room? An experimental analysis of information avoidance in a market setup”. In: *Ecological Economics* 169 (2020), p. 106479.
- [36] Mueller, Andreas, Spinnewijn, Johannes, and Topa, Giorgio. “Job seekers’ perceptions and employment prospects: Heterogeneity, duration dependence, and bias”. In: *American Economic Review* 111.1 (2021), pp. 324–63.
- [37] Oster, Emily, Shoulson, Ira, and Dorsey, E Ray. “Optimal expectations and limited medical testing: evidence from Huntington disease”. In: *American Economic Review* 103.2 (2013), pp. 804–830.
- [38] Rosholm, Michael and Toomet, Ott. “A search model of discouragement”. In: *Available at SSRN 742331* (2005).
- [39] Spinnewijn, Johannes. “Unemployed but optimistic: Optimal insurance design with biased beliefs”. In: *Journal of the European Economic Association* 13.1 (2015), pp. 130–167.
- [40] Statistics South Africa. *Quarterly Labour Force Survey: Quarter 2: 2022*. Statistical Report. Mar. 29, 2022. URL: <https://www.statssa.gov.za/?p=15685> (visited on 12/03/2022).
- [41] Tversky, Amos and Kahneman, Daniel. “Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty.” In: *Science* 185.4157 (1974), pp. 1124–1131.

- [42] Wheeler, Laurel, Garlick, Robert, Johnson, Eric, Shaw, Patrick, and Gargano, Marissa. “LinkedIn (to) job opportunities: Experimental evidence from job readiness training”. In: *American Economic Journal: Applied Economics* 14.2 (2022), pp. 101–25.
- [43] Zizzamia, Rocco. “Is employment a panacea for poverty? A mixed-methods investigation of employment decisions in South Africa”. In: *World Development* 130 (2020), p. 104938.
- [44] Zizzamia, Rocco and Ranchhod, Vimal. *Measuring employment volatility in South Africa using NIDS: 2008-2017*. Working Paper 246. Southern Africa Labour and Development Research Unit, 2019.

Appendices

A Top-up rounds

After all in-person sessions had been held, subjects were invited to participate in an additional ‘top-up’ session. This top-up session was made up of the two treatment conditions *SearchNoRej* and *SearchMoreRej* (ordering randomised).²⁴ The *SearchMoreRej* condition was introduced in the final two in-person sessions as well as in a top-up session for those who had participated in the first eight in-person sessions.

The *SearchNoRej* condition had not been included in the original in-person sessions, or in the pre-analysis plan (see Appendix G). During the experiment I became concerned that I would not be able to separate the rejection effect from potential framing effects, which motivated the design of the *SearchNoRej* condition. Like *SearchMoreRej*, this condition was introduced in the final two in-person sessions as well as in the top-up session.

In total, 122 out of a potential 203 participants completed the *SearchNoRej* or *SearchMoreRej* treatments in either the top-up session or in-person sessions.²⁵ There may be concerns that differential take-up of the top-up sessions may introduce bias into the comparisons between treatments implemented in the top-up sessions and those in the main sessions. Overall, apart from some gender imbalance, I do not find evidence for substantial imbalance on observables between these samples, as reported below.²⁶

²⁴*SearchMoreRej* had been included in the in-person sessions, but a coding error unfortunately invalidated results.

²⁵Participants in the top-up session were eligible for incentives as for the in-person sessions, but did not receive a show-up fee.

²⁶The gender imbalance is not a major concern. As reported in Table H.1 in Appendix H, gender does not appear to be correlated with search behaviour in the game.

Table A.1: Balance between sample that completed the top-up session and those that did not

Variable	(1)		(2)		T-test Difference (1)-(2)
	N	Mean/SE	N	Mean/SE	
Black African	82	0.793 (0.045)	121	0.826 (0.035)	-0.034
Female	82	0.476 (0.055)	121	0.669 (0.043)	-0.194***
Age (yrs)	82	23.537 (0.329)	121	23.347 (0.269)	0.189
Some tertiary education	82	0.671 (0.052)	121	0.669 (0.043)	0.001
Employment experience (mnths)	82	12.159 (1.413)	121	9.983 (1.179)	2.175
Currently employed	82	0.744 (0.048)	121	0.802 (0.036)	-0.058
Current earnings (ZAR)	42	3383.381 (152.287)	77	3464.987 (111.055)	-81.606

Notes: The value displayed for t-tests are the differences in the means across the groups. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level.

B Full list of treatment conditions

In the table below I describe the full list of treatment conditions that were implemented during the experimental sessions run in 2022. The first four are those described in Table 4.1 and are those which provide the data that form the empirical basis of this paper. In addition to these four treatment conditions, I also implemented two other treatment conditions, with varying sample sizes. In these four additional treatment conditions, I intended to explore whether the provision of additional information along with information on rejection/acceptance might affect individuals' behaviour in the game. In all cases, there was a monetary incentive to *search*.

In *PeerEffortInfo*, when a player received the rejection message, she also received information on the performance of her peers (those playing the game in the same session) at the real effort task, and on whether her performance was above the peer-group average,

or below the peer-group average. This was intended to provide a way to study whether information on peer performance would be motivating or demotivating, and whether this effect might depend on whether one is above or below the peer average. In *PeerOutcomeInfo*, when a player received the rejection message, she also received information on the *outcomes* of her peers (those playing the game in the same session) – i.e. out of x number of players, how many received an offer in that round. This was intended to provide a way to study whether information on peer *outcomes* would be motivating or demotivating for those who are rejected in a given round.

Of the 6 treatment conditions listed below, *Stay*, *Search*, *SearchMoreRej*, *PeerEffortInfo* and *PeerOutcomeInfo* were pre-specified in the pre-analysis plan (see Appendix G). During the period in which the experiment was run *SearchNoRej* was added to provide a means of testing for robustness to framing effects.

However, data from *PeerEffortInfo* and *PeerOutcomeInfo* are not included in the main paper for two main reasons: First, I felt that the question of the relationship between peer-information and behavioural responses to rejection was a distraction from the main focus of the paper, which is on the identification of a rejection effect. Secondly, I was not satisfied with the implementation of these treatment arms: While running the sessions, it appeared that participants quickly scanned the outcome information to determine whether or not they had been rejected. Beyond this, it did not appear that they paid any attention to the informational content of the rejection messages. Partly, this may have been because of the time pressure in the game, and partly because individuals were interested primarily in the outcome and not much else beyond this. The behaviour in these ‘information’ arms was largely similar to that in the *Search* condition, though given the concerns raised above, it is difficult to interpret the null effect of information provision and so results are omitted.

Table B.1: All treatments

Treatment	Payoff: finding job (ZAR)	Payoff: keeping job (ZAR)	Rejection signals when searching?	Extra information when rejected?	Round-by-round rejection rate (median effort)	n. rounds	Players
1. Stay	50	55	YES	No	96%	8	183
2. Search	55	R0	YES	No	96%	8	179
3. SearchNoRej	55	50	NO	No	96%	8	109
4. SearchMoreRej	55	50	YES	No	98%	16	109
5. PeerEffortInfo	55	50	YES	On performance relative to peers	96%	8	60
6. PeerOutcomeInfo	55	50	YES	On offers recieved by searching peers	96%	8	147

Notes: The round-by-round rejection rate is, for illustrative purposes, set at the (approximate) observed median effort (12/30 sliders completed). The probability of receiving an offer in any given round is determined by the player’s performance in the effort task in that round, and ranges from 0% to 10% in the 8-round games, and 0% to 5% in the 16 round game. At median effort, the *cumulative* probability of receiving an offer over the course of the game is approximately equal in Treatments 1-3 (27.9%) and Treatment 4 (27.6%).

In T5 T6 T7 and T8, players received “extra” information if they were rejected: In the *PeerEffortInfo* treatment players were told whether they had performed better or worse than their peers in the real effort task. In the *PeerOutcomeInfo* treatment players were told how many other players in their “room” (a group of between 8 and 12 players, grouped together in each session) had received an offer in the session.

C Comprehension tests

Players were required to complete the following comprehension test before beginning the *session*. As shown in the image, if players answered incorrectly, they were given the chance to correct their answers. After the comprehensions test I went over the questions, explaining why particular answers were correct and incorrect.

Figure 7: Main comprehension check

Comprehension Test

Please answer the following questions to make sure that you understood the instructions.

1) What determines whether you get offered a job or not?

- How many of the slider tasks you complete in a given round
- Luck
- Both answers

2) Which of these statements is true?

- If you search for a job, you will only find out at the end of the whole game whether you are offered a job.
- If you search for a job, you are told whether you were offered a job as soon as you finish the slider task.

3) Which of these statements is true?

- If you search for a job and get a job offer, to keep it you need to keep up the effort you put in before you got the job offer.
- If you search for a job and get a job offer, you can relax for the rest of the game - your airtime is guaranteed!

4) A player completes 15 out of 30 sliders in a round. How likely is she to get a job offer (whether new or current) in that round?

Hint: Remember, 3 sliders = 1/100 increase in your chances of getting an offer.

- 15/100 chance of getting a job offer
- 5/100 chance of getting a job offer
- 45/100 chance of getting a job offer

Next Page

Please fix the errors.

Players were required to complete versions of the following shortened comprehension test before beginning each *game*.

Figure 8: Supplementary comprehension check

Comprehension Test

Please answer the following questions to make sure that you understood the instructions.

1) Imagine you choose to search and complete 15 sliders, and the person next to you decides to focus on her current job and completes 15 sliders. Which of the following statements is correct?

- The person next to you who decided to focus on her current job is more likely to get an offer.
- You both stand the same chances of getting a job offer.
- You (who decided to search) are more likely to get an offer.

2) In this game do you stand to earn more if you search or if you focus on your current job?

- You can earn more if you find a new job
- You can earn more if keep your current job

Next Page

Please fix the errors.

D The effects of increasing the frequency of rejection

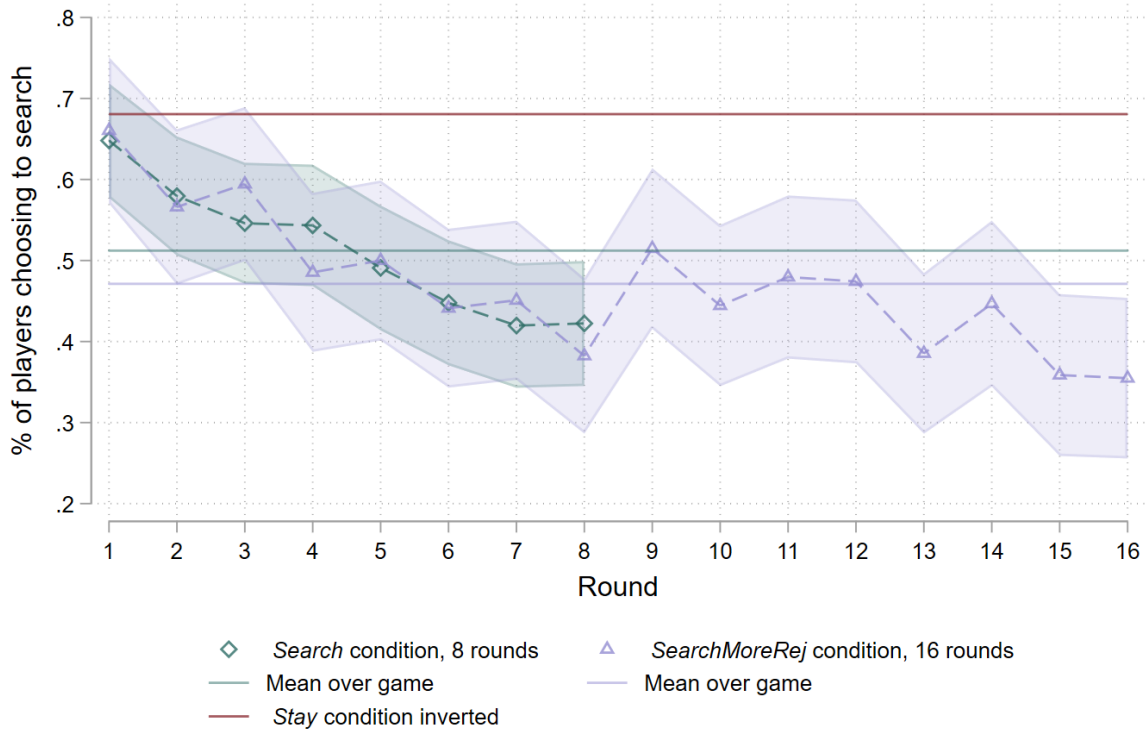
D.1 Design

A fourth treatment, *SearchMoreRej*, was intended to investigate non-linearities in the effects of repeated rejection. It is identical to *Search* except that the number of rounds is doubled and thus there is an increase from eight potential rejections to 16 potential rejections. Importantly, I hold the cumulative probability of receiving an offer roughly constant between *Search* (27.9 percent) and *SearchMoreRej* (27.6 percent). While the cumulative probability of receiving a payoff at the end of the game remains approximately unchanged, the likelihood of being rejected after any given round exactly doubles.

It is unclear *ex ante* whether the marginal psychological cost of the first rejection is higher than the marginal psychological cost of a rejection received after several prior rejections. One could imagine either that individuals get a ‘thick skin’ after repeated rejection and become habituated, or that they get increasingly ‘worn down’. Optimistic luck-based interpretations of rejection may be more difficult to sustain after repeated rejections, for instance. On the other hand, there is also likely heterogeneity across individuals in terms of their rejection aversion, with the result that highly rejection averse individuals will self-select out of search earlier than rejection insensitive individuals.

If the marginal effect of rejection on discouraging search is *decreasing* in exposure to rejection, we would expect the differences in search rates between *SearchNoRej* and *Search* to be greater than the differences in search rates between *Search* and *SearchMoreRej*. If the marginal effect of rejection on discouraging search is *increasing* in exposure to rejection, we would expect the opposite.

Figure 9: Search rates in *Search* vs *SearchNoRej*



D.1.1 Treatment 4: *SearchMoreRej*

Comparisons between the *SearchMoreRej* and *Search* treatments allow us to determine whether the marginal effect of rejection on discouraging search is *decreasing* or *increasing* in exposure to rejection. That is, is discouragement sustained after the first eight rejections or does it flatten out?

As in *Search*, in this treatment subjects have a monetary incentive *to search* (see Table 4.1). The key characteristic of the *SearchMoreRej* condition is that the number of rounds

is doubled and thus there is an increase from eight to 16 potential rejections. I hold the cumulative probability of receiving an offer constant between *Search* (27.9 percent) and *SearchMoreRej* (27.6 percent). Thus, while the cumulative probability of receiving a payoff at the end of the game remains approximately unchanged, the likelihood of being rejected in any given round doubles.

Figure 9 compares search rates in the *Search* and *SearchMoreRej* conditions. Perhaps surprisingly, average search rates in *Search* and *SearchMoreRej* are very similar, and as we shall see in Table D.1 are statistically indistinguishable. Analysing trends, we see that search rates in *SearchMoreRej* track search rates in *Search* very closely for the first eight rounds as we should expect, but then show no evidence of a further decline.²⁷ This suggests that the discouragement effects of rejection are steepest in early rounds, after which point search rates flatten out – perhaps due to the more rejection sensitive subjects selecting out of search.

In Table D.1 I test whether exposing subjects to additional rejection (while keeping cumulative probabilities of receiving an offer constant) further discourages job search. Results suggest that there is a small but insignificant effect of additional rejection on job search. Considering these results alongside those presented in Table 5.2 suggests that the rejection effect on job search is steepest with the first rejections, after which point it flattens out.

²⁷Search rates in the *SearchMoreRej* condition are noisier than in the *Search* condition due to a smaller sample size (see Appendix B), perhaps explaining the counter-intuitive apparent (though statistically insignificant) increase in search rates in the rounds immediately following round eight.

Table D.1: Treatment effects of rejection on search rates, comparing *Search* to *SearchMoreRej* conditions

	Full sample	Excluding job-finders
Search vs SearchMoreRej	0.025 (0.043)	0.067 (0.046)
Constant	0.467*** (0.021)	0.439*** (0.022)
Mean search rate for SearchMoreRej	0.48	0.48
% change, Search vs SearchMoreRej	5.30	14.03
Number of observations	3347	3056
Number of players	180	169
Indiv. Fixed Effects	Yes	Yes
Past outcomes in session	Yes	Yes
Excludes job-finders	No	Yes

Notes: * $p < .1$; ** $p < .05$; *** $p < .01$. Standard errors, clustered at the subject level, are in parentheses. The outcome variable in both specifications is a binary variable equal to 1 if a subject decides to ‘search’ in a given round, and equal to 0 if a subject decides to ‘stay’. The independent variable of interest, *Search vs SearchMoreRej*, is defined as a dummy equal to one if the individual was in the treatment condition of interest (*Search*) and equal to zero if the individual was in the comparison treatment condition (*SearchMoreRej*). Coefficients can be interpreted as the treatment effect on search rates of being in the *Search* condition relative to the *SearchMoreRej* condition. The first column reports results for the full sample, while the second column reports results only for the subsample of subjects who did not search and receive a job offer in either condition under consideration.

E Robustness to ordering effects

The order of treatments Search and Stay was randomised across experimental sessions so as to eliminate, in expectation, the role of ordering effects on outcomes. However, because the number of rounds was small, this randomisation may not have netted out ordering effects in practice.

Ordering effects might be concerning since behaviour in one game might be influenced by outcomes in a previous game. For instance, if a player chose to ‘stay’ in one game, and at the end was informed of a positive outcome, they might be more likely to ‘stay’ in the following game. Unfortunately, because *SearchNoRej* and *SearchMoreRej* were implemented later in an add-on session, ordering was not randomly varied between *Search*

and *Stay*, and *SearchMoreRej* and *SearchNoRej*. However, in Table E.1 I report results for a *Search* and *Stay*, the ordering of which was randomised, so as to test for whether results are sensitive to ordering effects.

Comparing the estimates in Table E.1 when ordering fixed effects are included to when they are excluded (Table 5.2) one can see that effects appear to be robust to ordering effects. Including ordering fixed effects seems to slightly inflate coefficients relative to effects estimated without ordering fixed effects, suggesting that the main results presented in the paper may be marginal under-estimates.

Table E.1: Treatment effects of rejection on search rates: Search vs Stay (inverted), with treatment ordering fixed effects

	With framing effects	Without framing effects
Search vs Stay (inverted)	-0.194*** (0.052)	-0.217*** (0.054)
Constant	0.592*** (0.056)	0.627*** (0.057)
Mean search rate for Stay (inverted)	0.69	0.69
% change, Search vs Stay (inverted)	-28.16	-31.55
Number of observations		
Number of players	3143	2944
Indiv. Fixed Effects	203	196
Past outcomes in session	Yes	Yes
Excludes job-finders	Yes	Yes
Restricted	Yes	No

Notes: * $p < .1$; ** $p < .05$; *** $p < .01$. Standard errors, clustered at the player level, are in parentheses. The outcome variable in both specifications is a binary variable equal to 1 if a player decides to search in a given round, and equal to 0 if a player decides to ‘focus on her current job’. The independent variable of interest (Search vs Stay (inverted)) is defined as a dummy equal to one if the individual was in the treatment condition of interest and equal to zero if the individual was in the comparison treatment condition. Coefficients can be interpreted as the treatment effect on search rates of being in the treatment condition of interest relative to the comparison condition. The first column reports results for the full sample, while the second column reports results only for the subsample of players who did not search and receive a job offer in the conditions under consideration.

F Direct test for framing effects

A comparison between *Stay (inverted)* and *SearchNoRej* in Table F.1 allows me to test directly for the presence of framing effects. We already have suggestive evidence that framing effects do not play a meaningful role in determining search behaviour, as evidenced by the roughly symmetrical search patterns in the SearchNoRej and Stay conditions in Figure 4 above. Table F.1 provides further evidence that framing effects are small and statistically insignificant, with null effects fairly tightly centered on zero.

Table F.1: Treatment effects on search rates, comparing *Stay (inverted)* to *SearchNoRej* conditions

	Full sample	Excluding job-finders
Stay (inverted) vs SearchNoRej	0.037 (0.052)	0.036 (0.052)
Constant	0.667*** (0.034)	0.670*** (0.034)
Mean search rate for SearchNoRej	0.48	0.48
% change, Stay (inverted) vs SearchNoRej	7.66	7.52
Number of observations	2542	2472
Number of players	203	196
Indiv. Fixed Effects	Yes	Yes
Past outcomes in session	Yes	Yes
Excludes job-finders	No	Yes

Notes: *p < .1; ** p < .05; *** p < .01. Standard errors, clustered at the player level, are in parentheses. The outcome variable in both specifications is a binary variable equal to 1 if a player decides to search in a given round, and equal to 0 if a player decides to ‘focus on her current job’. The independent variable of interest, *Stay (inverted) vs SearchNoRej*, is defined as a dummy equal to one if the individual was in the treatment condition of interest (*Stay (inverted)*) and equal to zero if the individual was in the comparison treatment condition (*SearchNoRej*). Coefficients can be interpreted as the treatment effect on search rates of being in the *Stay (inverted)* condition relative to the *SearchNoRej* condition. The first column reports results for the full sample, while the second column reports results only for the subsample of players who did not search and receive a job offer in either condition under consideration.

G Pre-analysis plan

The experiment was pre-registered on the AEA RCT registry, and can be found [here](#).

Of the four treatment conditions which form the empirical basis of this paper (listed in Table 4.1), three (*Stay*, *Search*, *SearchMoreRej*) were pre-specified. One treatment condition analysed in the main paper that was not pre-specified, for reasons described in Appendix B. There were also two treatment conditions listed in the pre-analysis plan (*PeerOutcomeInfo*) which were not analysed in the main paper, again for reasons described in Appendix B.

H Correlates of search behaviour and discouragement

The survey data collected during the experiment allows me to explore some of the demographic, labour market and psychological correlates of search behaviour in the laboratory game. Table H.1 reports multivariate regression coefficients of a range of demographic, labour market and psychological variables on two outcome variables representing the likelihood of being ‘discouraged’ by a prior rejection. The first outcome variable is defined as the proportion of all rounds *following a rejection in the prior round* in which an individual chose to ‘stay’.²⁸ For instance, suppose a particular subject was asked to make a decision between ‘searching’ and ‘staying’ on 20 occasions in the game, 10 of which immediately followed a rejection in the prior round. If the subject chose *not* to search on six out of those 10 post-rejection occasions, her value for the outcome variable would be 0.6. A higher value therefore represents more ‘discouragement’. The second outcome variable is simply the first outcome variable, weighted by the amount that subjects searched overall, this way taking into account the *duration* of the discouragement effect. If a subject is discouraged by rejection, but then switches back to searching quickly, the ‘discouragement index’ will be down-weighted heavily relative to the first outcome variable. If, on the other hand, a subject is discouraged by rejection and never returns to searching, the ‘discouragement index’ will be larger than the first outcome variable. This variable provides a more complete perspective on discouragement in the game, but has the disadvantage of being somewhat less intuitive to explain. For completeness, I include both.

²⁸Observations in the *Stay* and *SearchNoRej* conditions are excluded, since pecuniary returns incentivised ‘staying’ in the former, and no rejections were received in the latter.

Explanatory variables include basic demographic characteristics (race, gender, age and education), psychological variables – such as whether an individual attributes success to her own merit or to luck, and self-reported motivation to search – as well as labour market variables representing the subject’s labour market behaviour in the real world – such as whether or not an individual is currently searching, and if employed, the individual’s expectation of being retained in her current job. Variable definitions are provided in the table notes.

I run a simple OLS regression of the ‘discouragement’ outcome variables on these explanatory variables. Coefficients are reported in Table H.1. The first specification (column 1) only includes the four demographic variables, the second specification adds the two psychological variables, and the third column includes the variable for whether or not the individual is searching for jobs in ‘real life’. The final column adds the control variable indicating whether or not respondents have an above or below median expectation of being retained in their current jobs. Because this variable is only asked of respondents who are employed in real life, this cuts the sample size by about 20%.

Compared to coloured respondents, black subjects appear to be approximately 14 percentage points more likely to stop searching following a rejection. However, this apparently large effect is somewhat diluted when we account for the fact that black subjects search more than coloured respondents overall, with the result that the coefficients on the ‘discouragement index’ are smaller and significant at a lower level. Older respondents appear to be more resilient to rejection, with lower rates of switching following a rejection. However, this effect disappears when we account for the total number of periods spent searching, suggesting that while older respondents are less likely to switch following a rejection, they are also less likely to come back to searching once they have switched. Gender and education do not appear to be correlated with discouragement in the lab game. This contrasts with other findings that suggest that there may be a correlation between gender and job-search persistence: for instance, in a field experiment Bapna et al. (2021) find that women are less likely to apply for further positions after having been exposed to rejection. On the other hand, my own previous work (Zizzamia, 2020) has suggested that young men in South Africa have a weaker attachment to the labour market and may therefore be more likely to be discouraged by disappointment in the waged labour market than women – who often have stronger family responsibilities.

Subjects who tend to attribute successes in their lives to their own merit – a concept

related to self-efficacy and locus of control – appear to be less likely to switch away from search following a rejection than subjects who subscribe to more luck-based explanations for success. However, this effect does not persist when we account for the number of search spells, with coefficients losing significance for the ‘discouragement index’. Subjects who claimed to struggle to motivate themselves to search for jobs in real life are, unsurprisingly, more likely to be discouraged by a rejection. Interestingly, coefficients increase in size and significance when taking into account the total amount subjects search – indicating that subjects who struggle with motivation are not only more likely to cease searching following a rejection, but on average also take longer to (or are less likely to) return to search.

While a bivariate regression of the discouragement variables on search behaviour in real life (‘Currently searching’) returns a negative and significant coefficient (not reported here), once we include controls this result disappears, as seen in columns 3 and 4 – suggesting that other psychological and demographic factors, correlated with search behaviour in the real world, explain the bivariate association. Finally, restricting the sample to those who were employed in the real world, I find that those with high expectations of being retained at their real jobs display less resilience to rejection in the laboratory game. While it would be unsurprising to observe that those who have higher expectations of retention in real life also search less in real life, the fact that we see this search behaviour carry over into the laboratory game suggests that there may be a deeper psychological explanation for over-optimism and active avoidance of rejection information.

Table H.1: Correlates of search behaviour and discouragement

	% of rejections followed by switch away from search				Discouragement Index			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Black African	0.132*** (0.039)	0.140*** (0.043)	0.140*** (0.044)	0.142*** (0.043)	0.092*** (0.033)	0.090** (0.035)	0.090** (0.035)	0.065* (0.038)
Female	-0.041 (0.035)	-0.047 (0.035)	-0.047 (0.035)	-0.019 (0.037)	-0.025 (0.030)	-0.029 (0.029)	-0.030 (0.029)	-0.005 (0.031)
Age	-0.016*** (0.006)	-0.014** (0.006)	-0.014** (0.006)	-0.012* (0.006)	-0.007 (0.006)	-0.005 (0.006)	-0.005 (0.006)	-0.005 (0.005)
Some tertiary training	0.041 (0.036)	0.041 (0.036)	0.041 (0.036)	0.014 (0.037)	0.024 (0.033)	0.026 (0.033)	0.026 (0.033)	0.020 (0.033)
Attributes success to own merit		-0.067** (0.033)	-0.067** (0.033)	-0.045 (0.035)		-0.046 (0.029)	-0.045 (0.029)	-0.002 (0.031)
High self-efficacy		0.010 (0.054)	0.000 (0.061)	-0.034 (0.070)		-0.017 (0.043)	-0.019 (0.047)	-0.044 (0.054)
Struggles with motivation to search for jobs (self-reported)		0.069** (0.034)	0.069* (0.035)	0.055 (0.038)		0.091*** (0.031)	0.095*** (0.033)	0.075** (0.036)
Currently searching			0.002 (0.048)	-0.003 (0.054)			0.032 (0.040)	0.016 (0.044)
Higher than average expectation of being retained				0.082** (0.034)				0.059* (0.030)
Constant	0.567*** (0.133)	0.536*** (0.137)	0.535*** (0.138)	0.435*** (0.142)	0.289** (0.125)	0.231* (0.127)	0.217* (0.128)	0.169 (0.128)
R ²	0.08	0.14	0.14	0.19	0.04	0.11	0.12	0.11
Number of observations	149	146	146	113	156	152	152	117

Notes: *p < .1; ** p < .05; *** p < .01. Robust standard errors are in parentheses. The *race variable* is a dummy equal to 1 if the subject is black and equal to zero if the subject is coloured (there were no whites in the sample). *Female* is a dummy equal to 1 if a respondent identifies as a woman, and zero if a respondent identifies as a man (there were no gender non-confirming individuals in the sample). *Age* is simply age in years. *Some tertiary education* is a dummy equal to 1 if a subject had completed any post-schooling qualification, and zero if the subject had not progressed beyond secondary school. All subjects had at least completed secondary schooling. *Attributes success to own merit* is a dummy equal to 1 if the subject is deemed to attribute instances of success in her life to her own merit. This variable is constructed as follows: Respondents were presented with the following two pairs of statements: 'I have a better chance of finding work than most of my peers. vs I basically have the same chance of finding a job as my peers - finding a job is mostly due to luck.' and 'The key to finding a job in South Africa is luck. vs The key to finding a job in South Africa is to not give up.'. Respondents were then asked to choose which between the pairs of statements they agreed with. If they agreed with the statements in which success was attributed to explanations centered on the individual, the 'success attribution' variable was coded as 1. If the subject agreed with the either of the statements in which success was attributed to explanations external to the individual, the 'success attribution' variable was coded as 0. *Self-efficacy* is measured using the New Generalised Self-efficacy Scale (G. Chen et al., 2001). *Struggles with motivation to search for jobs* is a dummy variable constructed as follows: Respondents were asked on a likert scale to indicate how strongly they agreed or disagreed with the following statement: 'I struggle to motivate myself to search for jobs'. Those who 'Agreed' or 'Strongly agreed' are coded as 1, and those who disagreed or were neutral were coded as zero. *Currently searching* is a dummy equal to 1 if a respondent was searching at the time of the survey, and zero otherwise. *Higher than expected expectation of being retained* is constructed as follows: Respondents who were employed at the time were asked how likely they believed, on a scale from 1-10, it would be for them to be retained in their current jobs. I used this variable to construct a dummy, with those above median expectations of being retained coded as 1, and those with below median expectations of being retained coded as zero.