

EXPLORATION & SELF-SELECTION: REVISITING ROY*

Barton H. Hamilton[†]

Washington University in Saint Louis, Olin Business School

Stephanie A. Heger[‡]

University of East Anglia, School of Economics

July 22, 2020

ABSTRACT: We study how selection shapes exploratory versus exploitative behavior. Using a laboratory experiment, we decompose Roy's (1951) theory of selection and study the effect of traits on innovative behavior when information type is assigned versus when it is selected. Consistent with theory, when information type is assigned, we find that (1) there are distinct behavioral patterns leading to earnings' disparities and (2) the returns to personality traits are ambiguous, and significantly depend on the type of information assigned. By contrast, when individuals self-select their information type, we show that they leverage their trait-based advantage, their information choice is optimal and, as predicted by Heckman & Honore (1990), the earnings' disparities disappear.

KEYWORDS: selection, exploration, exploitation, traits

*We received helpful suggestions from seminar participants at Bocconi University, the University of Sydney and Washington University in Saint Louis, as well as comments and suggestions from Nicholas Papageorge, Filippo Massari, and Agnieszka Tymula on an earlier draft.

[†]Olin Business School, One Brookings Drive, Saint Louis, MO 63130. Email: hamiltonb@wustl.edu

[‡][Corresponding author] Level 5, Social Sciences Building, The University of Sydney, Sydney, NSW 2006. Email: stephanie.heger@sydney.edu.au

1 Introduction

A multi-armed bandit is a statistical decision model that represents the trade-off between acquiring more information about the underlying problem while simultaneously maximizing a stream of pay-offs (Robbins, 1952)—exploration versus exploitation (March, 1991). In economics, applications of bandit problems include search in labor markets (Jovanovic, 1979; Miller, 1984) and innovation and entrepreneurship (Weitzman, 1979; Manso, 2011; Ederer and Manso, 2013; Herz et al., 2014),¹ while in political science bandit problems have been used to model policy experimentation (Hirsch, 2016). Other disciplines, including psychology, management and neuroscience have focused on the *type* of people or organizations that succeed in balancing the exploration-exploitative trade-off inherent in multi-armed bandit problems. For example, psychologists refer to individuals who are better able to cognitively manage the contradictory goals of exploring new ideas and recognizing the benefits of exploiting the current opportunity as having a paradoxical frame (Amabile, 1983; Smith and Tushman, 2005). Similarly in the management and strategy literature, Tushman et al. (1996) refers to firms that are able to manage the exploitation-exploration trade-off as having organizational ambidexterity and Smith and Tushman (2005) suggest that individual managers who can engage in paradoxical thinking are a key component for organizational ambidexterity. There are also biological underpinnings that support the idea that exploration-exploitation problems involve competing cognitive processes (Aston-Jones and Cohen, 2005) and that different areas of the brain are involved in exploration versus exploitation (Daw et al., 2006).

In this paper, we ask how selection based on personality affects exploratory versus exploitative behavior. Recent experimental evidence suggests that information acquisition or exploration depends on individual personality (Herz et al., 2014; Fréchet et al., 2017). Thus, the self-selection of information, rather than personality itself could lead individuals, to different outcomes of the game. Roy (1951)’s model of occupation choice provides a framework for understanding the implications of the bias that can occur when individuals self-select.² The underlying difficulty in these questions is that the researcher only observes behavior and earnings conditional on self-selection, rather than for the entire population. We purposefully circumvent this problem by studying selection in a multi-armed bandit problem. We are not interested in studying which specific traits lead to more successful balancing of exploration and exploitation, but instead, we conjecture and show that individual traits play a significant role in the selection of information in exploratory versus exploitative endeavors.

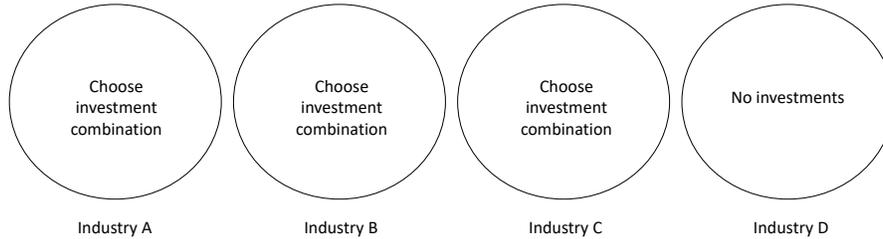
The main task in our experiment, the Industry Game, is adapted from Ederer and Manso (2013) and Herz et al. (2014), which is a multi-armed bandit problem.³ We add significant context to the multi-armed bandit to ease subject comprehension (Alekshev et al., 2017) and subjects in the

¹See also Bergemann and Valimaki (2006) for a review.

²Roy models have been used in a variety of contexts to better understand how the interaction of self-selection and individual characteristics drive different earnings’ patterns, including immigration (Borjas, 1987), college attendance (Willis and Rosen, 1979), and labor force participation (Gronau, 1974; Heckman, 1974).

³Ederer and Manso (2013) calculate that there are 6,181,806 combinations in their task. We have significantly reduced the dimensions of the game to 1,591,651 unique combinations or “arms”.

FIGURE 1: THE INDUSTRY GAME AS A MULTI-ARMED BANDIT



This figure provides an overview of the main task, the Industry Game, used in the experiment. The Industry Game is a nested multi-armed bandit. Each Industry has an unknown fixed cost drawn from a uniform distribution and within each Industry A, B and C, subject choose how to invest their endowment. In each Industry A, B C, there are 530,550 unique investment combinations. Industry D also has an unknown fixed cost, but there are no investment opportunities in Industry D and thus always pays out the Endowment minus the fixed cost. Thus, there are a total of 1,591,651 “arms” of the bandit. The two information sources, Cost Information and Investment Information, provide distinct types of feedback during the game. Cost Information provides feedback on the fixed cost, thus giving subjects a way of comparing across Industries or “nests” of bandits. Investment Information provides feedback on how change your investment strategy to increase your profit within the chosen Industry.

Industry Game take on the role of a manager in which they must decide which Industry to enter and how to invest their money across the Industry’s three products.⁴ The objective in each of the 20 rounds of the Industry Game is to maximize earnings. In each Industry, there is an unknown optimal product mix that maximizes the subject’s investment in the Industry and an unknown Industry-specific fixed cost. Thus, to maximize earnings subjects must decide when to explore new Industries or investment strategies and when to exploit (or fine-tune) their current strategy.

We employ a between-subject design where our main treatment manipulation is information in the Industry Game. Subjects are randomly assigned to one of four treatments: No Information, Investment Information, Cost Information or Information Choice (i.e., subjects choose the type of information they prefer to receive). The No Information treatment provides subjects only with earnings feedback after each round. By contrast, in the Investment Information and Cost Information treatments, we randomly assign subjects to receive either Investment Information *or* Cost Information, in addition to the earnings feedback. Investment Information consists of an unbiased signal about the optimal industry-specific investment level relative to their current investment

⁴We thus join a small, but notable literature examining innovation and creativity in the laboratory Charness and Grieco (2018); Ederer and Manso (2013); Herz et al. (2014); Meloso et al. (2009). Of these studies, only Herz et al. (2014) relates individual traits to innovative behavior and find that optimism is related to increased innovative behavior, while overconfidence is related to less innovation.

strategy (this is equivalent to the feedback in Ederer and Manso (2013) and Herz et al. (2014)) and allows subjects to exploit investment strategies *within* industries. By contrast, Cost Information consists of an unbiased signal about the subject’s industry-specific fixed cost and provides information that allows subjects to compare *across* industries. In the fourth treatment, the Information Choice treatment, subjects *choose* the type of information that want to receive: no information (Control), Investment Information, or Cost Information. Importantly, the first three treatments *assigns* subjects to an information sector (i.e., No Choice treatments), while the fourth treatment allows subjects to *self-select* into their preferred information sector. In addition to the Industry Game, we elicit Big Five personality traits (Costa and McCrae, 1985),⁵ locus of control Rotter (1971),⁶ risk preferences and cognitive ability Raven and Court (1998), thus also contributing to a growing literature of non-cognitive skills on economic outcomes (see Almlund et al. (2011) for an overview of this literature).

Our pattern of findings closely mirrors the predictions put forth by Roy (1951). In particular, there are significant exploration and earnings’ disparities when information is randomly assigned, but these disappear when subjects are able to leverage their comparative advantage and self-select into their preferred information sector (Heckman and Honore, 1990). For example, subjects randomly *assigned* to Cost Information explore more and earn significantly less than subjects randomly *assigned* to Investment Information. However, these differences disappear in the Information Choice treatment when subjects *self-select* either Investment or Cost Information. Roy’s prediction stems from the idea that when selection occurs, individuals are able to leverage their comparative advantage. This is exactly what we find: rather than a single set of traits or information type being universally advantageous, certain traits are assets when assigned Cost Information, but liabilities when assigned Investment Information and when subjects can *select* their Information type they predictably leverage their trait-based comparative advantage.⁷

In particular, we find that extraversion and risk tolerance are assets for subjects *assigned* to Cost Information, but liabilities for subjects *assigned* to Investment Information. These are the only two traits that have an opposite effect on earnings in the No Selection treatments. We find that Cost Information causes more exploration, however, we show that risk tolerance tempers subjects’ exploration when receiving Cost Information because they are more tolerant to variations in the expected fixed cost of the Industry (as inferred by the received signals) and thus are more willing to remain in that Industry and fine-tune their investment strategy than those with lower levels of risk tolerance. Extraverts, by definition, are inclined towards over-activity and we find that

⁵Recent literature links the Big Five to a host of factors that may affect labor market outcomes (Barrick and Mount, 1991; Caliendo et al., 2011; Fletcher, 2013; Hamilton et al., 2014; Cubel et al., 2016).

⁶Rotter’s External-Internal Locus of Control is designed to determine the extent to which an individual views his life as under his control. Individuals with an internal locus of control view their life as under their direct control and influence, a trait that is linked to need for high achievement and a preference for autonomy (McClelland, 1965) and subsequently to a preference for entrepreneurship (Brandstätter, 1997; Caliendo et al., 2011; Evans and Leighton, 1989).

⁷Similarly, Lundberg (2013) finds that personality traits interact with socioeconomic status such that Conscientiousness was associated with better educational outcomes for advantaged males, whereas Openness was associated with better outcomes for disadvantaged males.

instead of reaping the benefits of exploitation caused by assigned Investment Information, they are overly-active investors—they follow the advice given by the Investment Information feedback but they also engage in superfluous activity that reduces the benefits of exploitation.

Importantly, in the Information Selection treatment, extraversion and risk tolerance are the *only* traits that predict information selection—subjects who are more extroverted and risk tolerant are significantly more likely to select Cost Information than Investment Information.⁸ Further, in the Information Selection treatment, we find that risk tolerance and extraversion have no significant impact in how subjects respond to their informational signals. Finally, we show that subjects select optimally—on average, subjects who choose Cost (Investment) Information earn more using Cost (Investment) Information than they would have had they chosen Investment (Cost) Information.

2 Experimental Design & Data

The experiments were run at the University of Sydney in May and October 2014. Our sample consists of 208 subjects recruited through ORSEE (Greiner, 2015) and the experiment was programmed using Z-Tree (Fischbacher, 2007). Sessions lasted approximately 90 minutes and the average earnings were approximately 33 AUD. During the experiment, subjects could earn money during an Industry Game (20 Rounds), a lottery task (45 lottery choices) and a cognitive test (answer up to 12 questions, earn \$5 per correct question). This means, there were 66 items (20+45+1) for which the subject could earn money. At the end of the experiment, we randomly choose one of these decisions for payment. Additionally, subjects completed unincentivized personality and locus of control assessments. See Supplementary Material D for the experimental instructions and screenshots.

2.1 The Industry Game

The Industry Game used in our experiment is a version of the Lemonade Stand Task in Ederer and Manso (2013) and the Ice Cream Stand Task in Herz et al. (2014).⁹ While there are small changes in the structure of the game, the main elements remain the same.

In the Industry game, subjects take on the role of a manager who must decide how to invest resources for 20 rounds. At the beginning of each round, each subject i is endowed with 100 Australian dollars (AUD) and must make two choices: first, the subject chooses which of four industries to operate (Industry A, Industry B, Industry C, or Industry D); second, the subject decides how to invest in his chosen industry. Each subject has an unknown industry-specific fixed cost drawn randomly from a uniform distribution between 50 and 100, which remains fixed throughout the 20 rounds of the Industry game, $f_{i,I} \sim U[50, 100] \forall I \in \{A, B, C, D\}$. The subject knows that if he enters Industry A, B, and C he will have to make a positive investment by allocating his

⁸Fréchette et al. (2017) also finds evidence that personality predicts information demand.

⁹The authors thank Florian Ederer and Holger Herz for generously sharing their Z-Tree programs.

endowment across three investment products, x , y and z . The subject does not have to invest the entire endowment; any endowment that is not invested is considered savings for that round, although subjects are informed that savings do not carry over between rounds. The profit function is defined so that within each Industry, there is a unique, profit-maximizing investment strategy, $(x_I^*, y_I^*, z_I^*) \forall I \in \{A, B, C\}$. Subjects do not know the exact profit function, but they do know that their earnings depend on the amount invested, the distance their investment is from this bliss point, and their industry-specific fixed cost.¹⁰ Alternatively, subjects can exercise an outside option and enter Industry D. Industry D differs from the other three Industries in that there are no investment decisions to be made and subject always earns 100 minus his Industry D fixed cost. After an investment decision is made, the subject learns his earnings for the round and then proceeds to the next round. Subjects are also told that the maximum they can earn is 150 AUD (i.e., invest the entire endowment at the bliss point, which earns the subject 200 AUD and have the minimum possible fixed cost, 50 AUD) and that there is limited liability so any negative profits result in a payoff of 0 AUD.

There are four treatments: the Control treatment, the Investment Information treatment, the Cost Information treatment, and the Information Selection treatment. In the Control treatment, subjects play the Industry Game, as described above, and receive profit feedback after every round. The other three treatments provide profit feedback in every round as well as an additional piece of information, to be described, after each of the first 10 rounds.

Investment Information Treatment In the Investment Information treatment, subjects receive an unbiased signal about their investment strategy. The computer randomly determines whether to give information about one of the three products and then provides feedback about whether the subject should increase, decrease or not change the investment level in that product. For example, if a subject has over-invested in product x and product x is randomly chosen by the computer, then his signal will be to decrease his investment in product x . This information is equivalent to the “customer feedback” in Ederer and Manso (2013) and Herz et al. (2014).

Cost Information Treatment In the Cost Information treatment, in addition to profit feedback, subjects also receive an unbiased signal about their industry-specific fixed cost. The information is relevant to the Industry in which they are operating. Thus, if the subject is operating in Industry A, then he receives information about the fixed cost only in Industry A. For example, if a subject’s fixed cost in Industry A is 62, then the computer will randomly draw a number, z , from $Z \sim U [50, 100]$. If z is greater than 62, then the subject will receive a signal that says his fixed cost is less than z .¹¹

Information Selection Treatment In the Information Selection treatment, subjects choose whether they prefer to receive Cost Information, Investment Information or No Information during

¹⁰Appendix Supplementary Material C.1 shows the Industry-specific bliss points and profit functions.

¹¹S2 formally describes the signals.

the first 10 rounds. Before the game begins, subjects are shown each type of information and then asked to choose a single type of information to receive throughout the first 10 rounds. This treatment is designed to explore whether certain types of individuals prefer one type of information over the other and whether personality indirectly affects innovation through information choice.

Rounds 1-10 are an information accumulation phase. Investment Information and Cost Information are quite different forms of feedback in the Industry Game and thus, upon reaching Round 11, subjects assigned to the Investment Information Treatment have accumulated significantly different types of knowledge than subjects in the Cost Information Treatment. Investment Information provides highly specific feedback with explicit advice about how to increase profits. The individual simply needs to follow the advice to increase or decrease an investment in a product and their profits will increase. On the other hand, Cost Information is sparser and does not contain explicit advice. Instead, individuals must make additional inferences to effectively use Cost Information. For example, Cost information tells individuals whether their cost is above or below a randomly drawn number from the same distribution. Based on this information, the individual must decide whether to stay or leave the Industry. The updated belief that an individual has about his expected fixed cost in an industry will vary with respect to the expected value of the fixed cost and the expected variance. As we show in Section 3, the variation in expected variance plays an important role on the type of individuals who are able to more successfully use Cost Information.

2.2 Risk preferences, cognitive and non-cognitive skills

After subjects completed the Industry Game, we elicited risk preferences, cognitive ability, and personality traits. During the experiment, the elicitation of personality was always the final task. During approximately half of our sessions, we elicited risk preferences before cognitive ability and switched the order for the other half. We conduct all four treatments of the Industry Game with both task orders.

Risk preferences We elicit risk preferences following Hey and Orme (1994). Subjects faced a series of 45 lottery pairs and were asked to choose which lottery in the pair they preferred. We then follow Andersen et al. (2014) and estimate risk preferences at the individual-level, assuming CRRA utility, via maximum likelihood.

Cognitive Skills We use the Raven’s Advanced Progressive Matrices test to measure cognitive ability (Raven and Court, 1998), an intelligence test that is designed to be culture-free since it does not rely on language or cultural references. The test consists of 12 diagrams with a missing piece and eight suggested answers to the missing piece. The subject’s task is to choose one of the eight suggested answers. During the experiment, subjects have 12 minutes to complete 12 questions without feedback. We measure their cognitive ability as the number of correct answers.

Personality Traits We use the Big Five Personality inventory to assess personality.¹² We measured the Big 5 using the 120 item short form developed by Johnson (2014).

We use Rotter’s External-Internal Locus of Control test to measure locus of control (Rotter, 1971). The test consists of 29 pairs of statements and subjects are asked to indicate which of the two statements are consistent with their own views. The contemporary scoring system, which is the opposite of Rotter’s original scoring rule, associates higher scores with a more internal locus of control.

2.3 Data

Table 1 presents summary statistics of our sample. Note that the sample size is 194, rather than 208, due to technical difficulties in a session in which data from the Industry Game was collected, but data from the risk elicitation, cognitive test, and personality surveys were lost. The Big Five personality test is designed so that the median score for each trait is 50, with a standard deviation of 10. Also consistent with other findings, the subjects in our experiment are weakly risk-averse, with an average estimated CRRA coefficient of .89. Half of our subjects are female and the average age is just under 23 years.

The Industry Game is designed to measure degrees of exploration, but can also distinguish between exploration and “successful innovation”. Throughout our analysis, our main outcome variables are (1) exploration, and (2) earnings.

Exploration Ederer and Manso (2013) and Herz et al. (2014) measure exploration as the subject’s average industry-specific standard deviation in investment strategies. This measure captures the variance in the subject’s investment strategies but does not capture the frequency with which the subject changes industries. A change in the industry is perhaps the biggest exploration since it requires an entirely new and unknown investment strategy and, in our setting, an unknown fixed cost. Our measure of exploration, the Exploration Index, captures the degree of change in investment strategies and industry switches into a single measure.¹³ The Exploration Index scores the subject’s industry choice and investment strategy by how similar it is to all previous investment choices within the industry and assigns a score based on its similarity to the most similar strategy previously used. This allows us to identify when a subject returns to a previously tried idea (even when that choice happened several rounds before). We normalize the index between 0 and 1, inclusive. If a subject exactly replicates a previously used industry-investment choice or enters

¹²The Big 5 include Extraversion, Openness, Conscientiousness, Agreeableness and Neuroticism. Extraversion is associated with high energy, assertiveness, and positive affect. Openness reflects the degree of intellectual curiosity, creativity and is associated with a preference for a variety. Conscientiousness is associated with a tendency to be organized, efficient, dependable, and self-disciplined. Agreeableness is associated with the tendency to seek compromise and cooperation. Neuroticism is associated with being emotionally unstable and a tendency to experience anxiety and anger.

¹³In the Supplementary Material, we show that we obtain qualitatively equivalent results using the measure of exploration proposed in Ederer and Manso (2013) and Herz et al. (2014).

TABLE 1: SUMMARY STATISTICS

	All	Control	Investment	Cost	Selection
Openness	46.03 (9.01)	43.41 (9.07)	46.91 (9.52)	45.88 (9.50)	46.80 (8.40)
Extraversion	49.09 (8.04)	49.80 (9.87)	47.43 (7.48)	51.54 (8.01)	50.32 (7.29)
Neuroticism	48.96 (7.61)	48.22 (7.91)	48.80 (8.30)	52.54 (7.91)	48.26 (6.80)
Conscientiousness	49.36 (8.77)	49.76 (9.92)	49.80 (7.45)	45.89 (8.49)	50.47 (8.80)
Agreeableness	48.18 (8.44)	47.97 (9.92)	49.48 (6.98)	45.40 (8.60)	48.79 (8.26)
Locus of Control	11.49 (3.95)	11.97 (3.74)	11.43 (4.44)	11.31 (4.48)	11.37 (3.55)
CRRA coefficient	.81 (.78)	.85 (.53)	.59 (.51)	.87 (.97)	.89 (.89)
Raven Score, Cognitive Ability	7.20 (2.36)	7.19 (2.19)	7.20 (2.87)	7.20 (1.81)	7.19 (2.39)
Female	.55 (.50)	.60 (.49)	.43 (.50)	.60 (.50)	.60 (.50)
Age	22.74 (3.85)	22.33 (3.68)	22.61 (3.47)	22.57 (2.66)	23.06 (4.54)
Observations	194	36	44	35	79

We were unable to estimate risk preferences for 8 subjects. See Table S1 for more detail on sample sizes.

Industry D, then his Exploration Index in this round is 0. When a subject enters an Industry for the first time, his Exploration Index is 1.

We obtain the Exploration Index for subject i in period j in the following way. Define $I_{i,j} \in \{A, B, C, D\}$ be the industry chosen by subject i in period j . Let $(x_{i,j}, y_{i,j}, z_{i,j})$ be a vector of subject i 's investment strategy in period j . Define the Exploration Index of subject i in period j as follows

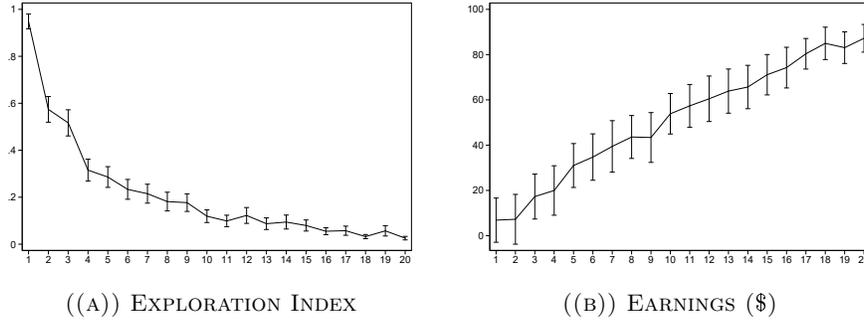
$$EI_{i,j} = \begin{cases} 0 & \text{if } I_{i,j} = D \\ 1 & \text{if } \forall j' < j \ I_{i,j'} \neq I_{i,j} \\ \kappa \times \min_{j' | I_{i,j'} = I_{i,j}} |x_{i,j} - x_{i,j'}| + |y_{i,j} - y_{i,j'}| + |z_{i,j} - z_{i,j'}| & \text{otherwise.} \end{cases} \quad (1)$$

where $\kappa = \frac{1}{200}$, which is the maximum deviation possible between two investment strategies, normalizes the Exploration Index so that it is between 0 and 1.¹⁴ The average Exploration Index with 95% confidence intervals for each of the 20 periods is shown in Figure 2(a) with 95% confidence bands.

Successful Innovation We also measure the degree to which subjects successfully innovate, which we measure in terms of money earned (see Figure 2(b)). Figure 2(b) shows the average earnings in each period. The trend shows that subjects perform better as the game unfolds.

¹⁴For example consider an investment strategy in period 1, $(x_{i,1}, y_{i,1}, z_{i,1}) = (100, 0, 0)$ and an investment strategy in period 2 of $(x_{i,2}, y_{i,2}, z_{i,2}) = (0, 100, 0)$ in Industry I . Then, the Exploration Index is given by $\frac{200}{1} \times \kappa = 1$.

FIGURE 2: OUTCOMES: EXPLORATION INDEX AND EARNINGS



Additional Control Variables In addition the control variables of interest (i.e., to personality and risk), we also include a set of control variables throughout our analysis. First, we exclude data from the first round of play since subjects make round 1 choices without any information and thus this choice is as good as random and only introduces noise. However, we do control for the pay-off the subject receives in round 1, since a “lucky” choice in round 1, and thus a lucky high pay-off, might influence how the subject plays the industry game. Second, we include fixed effects for cognitive ability (i.e., the number of correctly answered questions from the Raven’s test), round of play in the industry game (2-20), age, year in school, and order of play (i.e., some sessions completed the Raven’s test before the risk and other sessions performed the tasks in the opposite order).

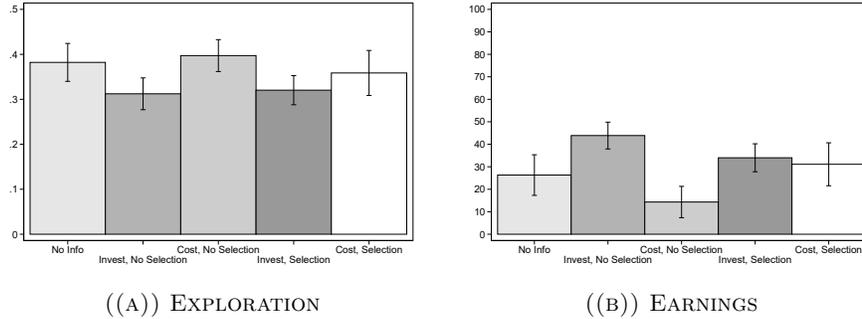
2.4 Effects of Information

Prior to exploring the role of traits, we first examine whether our main treatment manipulation (i.e., information types) results in differential outcomes for innovation and earnings.¹⁵ Figure 3 shows the average outcomes for exploration and earnings during the first 10 rounds by treatment (i.e., (1) No Information; (2) Investment Information; (3) Cost Information; (4) Information Selection-split by selection). We make three important observations. First, subjects assigned to the Cost Information treatment explore significantly more during the first 10 rounds than subjects assigned to the Investment Information treatment. This means, that at the conclusion of the first 10 rounds subjects assigned to receive Cost Information have experienced a wider breadth of investment strategies and industry choice combinations due to their greater propensity for exploration than subjects assigned to receive Investment Information.¹⁶ Due to their lower propensity to explore, subjects assigned to the Investment Information treatment have more finely-tuned and specialized knowledge. We conclude that two types of innovators emerge: Cost Information generates

¹⁵In Table S2, we present evidence that shows that subjects effectively use the information they receive by changing industries or adjusting their investment strategy.

¹⁶In Table S5 we follow the measurement of exploration in Ederer and Manso (2013) and Herz et al. (2014) and show the average standard deviation in investment strategies is significantly greater for subjects in the No Selection Cost Treatment than in the No Selection Investment Treatment.

FIGURE 3: AVERAGE TREATMENT EFFECTS



innovators that look like “Jack of All Trades” while Investment Information results in innovators that behave like “Specialists”.

Second, the fine-tuning strategy of the subjects assigned to Investment Information appears to be advantageous; random assignment to Investment Information, compared to Cost Information, leads to significantly more earnings. However, our third observation rules out the idea that Investment information is necessarily a superior form of information. Third, as predicted by Roy (1951), the innovation and earnings disparities *disappear* when subjects have the opportunity to select their preferred type of information.

As mentioned in Section 2.1, the differences in the type of information provided by the Investment information and the Cost Information treatments are significant. Investment information provides specific and explicit advice about how to increase profits, while Cost Information does not and instead requires subjects to make their own inferences about how to use the information to increase profits. Given these differences, it is not surprising that individuals may have preferences over information type. For example, and as we will see in Section 3, individuals who are more activity- and excitement-oriented (i.e., more extraverted) over-explore and subsequently achieve significantly lower earnings when they are assigned to Investment Information, but earn significantly more when assigned to Cost Information.

2.5 The Role of Traits

We now examine the role of traits on innovative behavior. To do so, we regress our two outcome measures—Exploration Index and Earnings—on a vector of individual traits and treatment dummies using data from the No Selection Treatments only (i.e., when information sector is exogenously assigned). In sum, we find that the Big Five personality traits are not jointly predictive of exploration or earnings and that there is no specific trait that plays a significant role.

TABLE 2: NO SELECTION TREATMENTS: THE ROLE OF INDIVIDUAL TRAITS

	Exploration	Earnings
Investment Info	-0.06*** (0.02)	17.27** (8.14)
Cost Info	-0.004 (0.02)	-9.47 (9.14)
Extraversion	0.0005 (0.001)	-0.37 (0.48)
Openness	0.0007 (0.001)	-0.52 (0.41)
Neuroticism	0.001 (0.001)	-0.13 (0.59)
Agreeableness	-0.001 (0.001)	0.64 (0.45)
Conscientiousness	-0.0006 (0.001)	0.04 (0.48)
Risk Tolerance	-0.008 (0.009)	7.09* (4.04)
Internal Locus of Control	-0.004** (0.002)	2.30*** (0.69)
Female	0.02 (0.02)	-14.30* (7.54)
Constant	0.55*** (0.16)	64.77 (63.32)
Observations	2074	2074
R^2	0.3	0.23
<i>F</i> -test		
Cost Info=Invest Info	9.19***	11.38**
Big Five traits	.65	.88
Controls		
Cognitive Skill FE	Y	Y
Round FE	Y	Y
Round 1 Pay-Off	Y	Y
Age & Year in School FE	Y	Y
Order FE	Y	Y

TABLE 3: OLS estimates. Robust Standard Errors clustered at the subject-level in parentheses and *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

2.6 Hypotheses

In Sections 2.4 and 2.5, we established two findings: (1) innovation and earnings’ disparities emerge when information is randomly assigned, but disappear when information is chosen; and (2) traits do not unambiguously drive innovative behavior. These two findings suggest that traits and information interact and that we may expect to find a pattern of predictable pattern of selection. To preview, our hypotheses and results are structured by decomposing a selection model (Roy, 1951). We hypothesize and show that (1) first, the returns to traits are information-dependent; (2) second, individuals’ demand for information is trait-based; and (3) third, individuals’ trait-based demand for information is optimal. Alternatively, in Supplementary Material B.1 we estimate a structural selection model and come to similar conclusions. We prefer the approach here because

it more clearly shows the nature of selection and how it relates to the question of understanding the “innovative personality”.

Hypothesis 1. *Information interacts with individual traits to drive innovation. The return to traits and information are interdependent.*

Our first hypothesis posits an interaction effect between traits and information. To test this hypothesis, we estimate equation 2 for subjects assigned to Investment Information and Cost Information, separately.

$$Y_{i,j} = \beta_0 + \beta_{\text{Traits}} \times \mathbf{X}_i + \beta_{\text{Controls}} \times \mathbf{Z}_i + \eta_{i,j} \quad (2)$$

If the interaction effects between traits and information are sufficiently strong, then, following Roy (1951), we expect that (1) information demand will be trait-based and (2) individuals optimally demand information. We turn to these hypotheses now.

Hypothesis 2. *Individuals will demand information that leverages their trait-based advantage. In particular, if a trait is an asset when assigned Investment Information, but a liability when assigned Cost Information, then an individual with this trait will be more likely to choose Investment Information.*

We test this information demand hypothesis using data from the Information Selection treatment and estimating the following probit regression

$$Pr[\text{Cost Information} = 1] = P_0 + \mathbf{P}_{\text{Traits}} \times \mathbf{X}_i + \varepsilon_i, \quad (3)$$

where the outcome variable takes a value of 1 if subject i chooses Cost Information and a value of 0 if the subject chooses Investment Information.

Our third hypothesis pushes the trait-based advantage further to better understand the nature of the selection problem. We hypothesize that individuals not only leverage their trait-based advantage through information demand but that they do so optimally; that is, on average, individuals could not have done better had they chosen a different type of information in the Information Selection treatment.

Hypothesis 3. *Individuals who chose Investment (Cost) Information could not have made more money choosing Cost (Investment) Information.*

To construct the counterfactual estimates of earnings and successful for subjects in the Information Selection Treatment, we use the estimates obtained from estimating equation 2 to predict the counterfactual outcomes. For subjects who chose Investment (Cost) Information, we use the estimated effects of individual traits from the average individual assigned to Cost (Investment) Information to predict what these subjects would have made if they had chosen the other type of information. We then construct four residual terms and test whether the residuals are consistent with subjects choosing optimally.

$$\begin{aligned}
E[\text{Dist To Optimum}_1|\text{Invest Info}=1] - E[\text{Dist To Optimum}_2|\text{Invest Info}=1] &< 0 \\
E[\text{Earnings}_1|\text{Invest Info}=1] - E[\text{Earnings}_2|\text{Invest Info}=1] &> 0 \\
E[\text{Dist To Optimum}_2|\text{Cost Info}=1] - E[\text{Dist To Optimum}_1|\text{Cost Info}=1] &< 0 \\
E[\text{Earnings}_2|\text{Cost Info}=1] - E[\text{Earnings}_1|\text{Cost Info}=1] &> 0
\end{aligned} \tag{4}$$

We estimate equation 4 by regressing (via OLS) the difference in the outcome variable in the chosen information sector with the predicted outcome variable in the alternative information section on a vector of individual traits and a constant. Thus, the constant represents the average difference in the residual, controlling for individual traits. A positive (negative) constant in the Earnings (Distance to Optimum) indicate that, on average, individuals perform better in their chosen information sector than they would have in the alternative.

3 Main Findings

In this section, we test each of the hypotheses described in the previous section. We begin with a statement of the result, followed by a brief discussion.

Result 1. *Individual traits interact with information to drive innovation. In particular, Extraversion and risk tolerance are assets when using to Cost Information, but liabilities when using Investment Information.*

In Table 4, we present the estimates from equation 2 to test Hypothesis 1. We find that increased Extraversion and risk tolerance is a liability when assigned Investment Information but an asset when assigned Cost Information. For example, a standard deviation increase in Extraversion leads to an average 12 dollar loss in earnings in Investment Information, but a 16 dollar gain in Cost Information. By contrast, Locus of Control, Neuroticism, and Agreeableness play similar roles in the exploration and earnings for both types of information.¹⁷

Why do more extraverted and more risk-tolerant individuals respond differently to Investment Information versus Cost Information? An underlying facet of the Extraversion trait is tendency towards activity and excitement-seeking (?). Thus, we may expect that the Investment Information treatment is too restrictive and does not allow for the type of self-directed activity favored by extraverts. In fact, this is exactly what the data shows—when receiving Investment Information, extraverts are equally likely to follow the advice given, but they *also* engage in additional exploration. For example, in the No Selection treatments, when extraverts receive a signal to “Increase your investment in x”, they are equally likely to increase their investment in x, but individuals scoring high on extraversion are significantly more likely to also change their investments in products y

¹⁷There are traits that play a significant role for one type of information and an insignificant role for the other type of information. We focus on those traits that have significant and opposite effects.

and z and change industries (see Table S3). Importantly, this heterogeneity in the responsiveness to signals disappears in the Selection treatments.

Individuals who are more risk-tolerant explore significantly less and significantly earn more when they receive Cost Information. Similar to the Extraverts' differential response to Investment information, we find that increasing risk-tolerance is associated with a differential response to Cost Information signals. Cost Information is useful because it allows a subject to update his belief about the expected value of the fixed cost in the current Industry as well as the expected variance and then decide whether to change industries or remain in his current industry given his updated beliefs. We find that when Cost Information is assigned (i.e., No Selection treatments), subjects with greater risk-tolerance are less responsive (i.e., less likely to change industries) to changes in expected variance of the expected fixed cost, holding constant the expectation of the fixed cost (see Table S4). This decrease in responsiveness allows more risk-tolerant subjects to decrease their exploration relative to others who receive Cost Information focus on exploiting their investment strategy *within* an Industry and subsequently increase earnings. Again and importantly, when subjects *select* Cost Information, we no longer find that risk tolerance is associated with less responsiveness to changes in expected variance of an Industry's fixed cost.

Result 2. *Individuals leverage their trait-based advantage when demanding Information. Increased Extraversion and risk tolerance is associated with a significantly increased likelihood of choosing Cost Information.*

Next, we turn to the Information Selection Treatment, where subjects self-select into receiving Cost Information or Investment Information after they have had a chance to learn about each type of information.¹⁸ Of the 79 subjects assigned to the Information Selection treatment, 52 chose Investment Information and 27 chose Cost Information. In Table 5, we present estimates from equation 3 and find that an increase in one standard deviation in Extraversion and Risk Tolerance is associated with 20 percentage point and 12 percentage point increase, respectively, in the likelihood of choosing Cost Information.¹⁹

Result 3. *Individuals optimally choose Information type; that is, individuals who choose Investment (Cost) Information earn more than they would have if they had chosen Cost (Investment) Information.*

Table 6 presents estimates from equation 4 to test whether individuals earn more in their chosen information sector than they would have if they had chosen the alternative information sector. We report the estimated mean residual calculated at the average of the covariates of personality, risk, locus of control and cognitive ability. Overall, subjects have higher earnings in the information regime they selected into than they would have in the alternate information regime.

¹⁸They also had the choice to choose No Information (i.e., the Control Treatment), but no subject made this choice.

¹⁹By contrast, Supplementary Material B.2 shows the effect of traits on innovation and earnings in the Information Selection treatment if we *ignore* their role in information demand.

TABLE 4: NO SELECTION TREATMENTS: EFFECT OF INDIVIDUAL TRAITS ON OUTCOMES, BY TREATMENT

	No Information		Investment Info Only		Cost Info Only	
	Exploration	Earnings	Exploration	Earnings	Exploration	Earnings
Extraversion	0.0000113 (0.002)	-0.01 (0.65)	0.004** (0.002)	-1.32*** (0.48)	-0.0007 (0.002)	1.63** (0.66)
Openness	0.001 (0.002)	-1.59* (0.94)	-0.0003 (0.0008)	0.08 (0.33)	-0.0009 (0.001)	-0.78** (0.38)
Neuroticism	0.0006 (0.002)	-0.77 (0.81)	0.003** (0.001)	-1.69*** (0.52)	0.004*** (0.002)	-1.26*** (0.47)
Agreeableness	0.001 (0.002)	-0.25 (0.68)	-0.002** (0.0009)	0.39 (0.52)	-0.006*** (0.002)	4.32*** (0.45)
Conscientiousness	-0.001 (0.002)	0.15 (0.8)	-0.0000669 (0.001)	-1.17* (0.61)	0.002 (0.002)	-0.53 (0.41)
Risk Tolerance	0.009 (0.03)	21.37** (10.57)	0.03 (0.02)	-24.92*** (7.05)	-0.02** (0.008)	18.03*** (2.21)
Internal Locus of Control	-0.002 (0.006)	-0.79 (1.81)	0.0002 (0.002)	-0.06 (0.59)	-0.001 (0.003)	0.91 (0.74)
Female	0.08*** (0.03)	-39.50*** (11.14)	0.03 (0.02)	-7.12 (8.78)	-0.02 (0.03)	24.27*** (8.69)
Constant	0.28 (0.23)	254.51*** (86.48)	0.24 (0.16)	356.09*** (78.38)	0.64** (0.27)	-315.68*** (87.78)
Observations	682	682	769	769	623	623
R^2	0.35	0.26	0.22	0.34	0.5	0.42
F -test						
Big Five traits			2.77**	3.70***	5.05***	34.22***
Controls						
Cognitive Skill FE	Y	Y	Y	Y	Y	Y
Round FE	Y	Y	Y	Y	Y	Y
Round 1 Pay-Off	Y	Y	Y	Y		
Age & Year in School	Y	Y	Y	Y	Y	Y
Order FE	Y	Y	Y	Y	Y	Y

4 Conclusion

Understanding *who* explores and exploits successfully has been a fascination among scholars dating back at least to Knight (1921), spurring literature in fields from economics and management to psychology and neuroscience.

In this paper, we step back from the traditional approach of studying exploration and exploitation and use a laboratory experiment. By doing so, we are able to directly study the interaction of individual traits and selection of information, in the form of information acquisition, that drives exploration-exploitation behavior. We find that there is no individual trait that unambiguously drives exploration or successful innovation, but rather that traits drive information demand and jointly determine innovative behavior. We find that individuals leverage their trait-based advantage and, when given the opportunity, optimally demand information.

TABLE 5: SELECTION TREATMENT: EFFECT OF INDIVIDUAL TRAITS ON INFORMATION DEMAND

	$Pr[\text{Cost Info}=1]$		
Risk Tolerance	0.14* (0.07)	0.13* (0.08)	0.16** (0.08)
Neuroticism	0.01 (0.01)	0.01 (0.01)	0.02 (0.01)
Conscientiousness	0.002 (0.008)	0.003 (0.008)	0.008 (0.009)
Openness	-0.007 (0.007)	-0.008 (0.007)	-0.006 (0.008)
Agreeableness	0.01 (0.009)	0.01 (0.009)	0.01 (0.009)
Extraversion	0.02** (0.01)	0.02** (0.01)	0.02* (0.01)
Internal Locus of Control	-0.002 (0.02)	-0.01 (0.02)	-0.002 (0.02)
Cognitive Ability	-0.002 (0.02)	-0.01 (0.02)	.
Female	0.1 (0.14)	0.05 (0.14)	0.02 (0.16)
Observations	76	76	73
Pseudo R^2	.	.	.
Controls			
Age, Year	No	Yes	Yes
Cognitive Ability FE	No	No	Yes
Task Order FE	No	Yes	Yes

Marginal effects from a probit regression. Robust Standard Errors in parentheses and *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

TABLE 6: COUNTERFACTUALS: TEST OF RESIDUALS

	Investment Info Only Earnings	Cost Info Only Earnings
Constant	15.77** (6.90)	19.97* (10.53)
Observations	980	540

OLS estimates. Robust Standard Errors clustered at the subject-level in parentheses and *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively.

References

- ALEKSEEV, A., G. CHARNESS AND U. GNEEZY, “Experimental methods: When and why contextual instructions are important,” *Journal of Economic Behavior & Organization* 134 (2017), 48–59.
- ALMLUND, M., A. L. DUCKWORTH, J. J. HECKMAN AND T. D. KAUTZ, “Personality psychology and economics,” Technical Report, National Bureau of Economic Research, 2011.
- AMABILE, T. M., “The social psychology of creativity: A componential conceptualization,” *Journal of personality and social psychology* 45 (1983), 357.
- ANDERSEN, S., G. W. HARRISON, M. I. LAU AND E. E. RUTSTRÖM, “Discounting behavior: A reconsideration,” *European Economic Review* 71 (2014), 15–33.
- ASTON-JONES, G. AND J. D. COHEN, “An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance,” *Annu. Rev. Neurosci.* 28 (2005), 403–450.
- BARRICK, M. R. AND M. K. MOUNT, “The big five personality dimensions and job performance: a meta-analysis,” *Personnel psychology* 44 (1991), 1–26.
- BERGEMANN, D. AND J. VALIMAKI, “Bandit problems,” (2006).
- BORJAS, G., “Self-selection and the earnings of immigrants,” *AMERICAN ECONOMIC REVIEW* 77 (1987), 531–53.
- BRANDSTÄTTER, H., “Becoming an entrepreneur—a question of personality structure?,” *Journal of economic psychology* 18 (1997), 157–177.
- CALIENDO, M., F. FOSSEN AND A. KRITIKOS, “Personality Characteristics and the Decision to Become and Stay Self-Employed,” (2011).
- CHARNESS, G. AND D. GRIECO, “Creativity and incentives,” *Journal of the European Economic Association* (2018).
- COSTA, P. T. AND R. R. MCCRAE, *The NEO personality inventory: Manual, form S and form R* (Psychological Assessment Resources, 1985).
- CUBEL, M., A. NUEVO-CHIQUERO, S. SANCHEZ-PAGES AND M. VIDAL-FERNANDEZ, “Do Personality Traits Affect Productivity? Evidence from the Laboratory,” *The Economic Journal* 126 (2016), 654–681.
- DAW, N. D., J. P. O’DOHERTY, P. DAYAN, B. SEYMOUR AND R. J. DOLAN, “Cortical substrates for exploratory decisions in humans,” *Nature* 441 (2006), 876–879.
- EDERER, F. AND G. MANSO, “Is Pay for Performance Detrimental to Innovation?,” *Management Science* (2013).

- EVANS, D. S. AND L. S. LEIGHTON, “Some empirical aspects of entrepreneurship,” *The American Economic Review* (1989), 519–535.
- FISCHBACHER, U., “z-Tree: Zurich Toolbox for Ready-Made Economic Experiments,” *Experimental Economics* 10 (2007), 171–178.
- FLETCHER, J. M., “The effects of personality traits on adult labor market outcomes: Evidence from siblings,” *Journal of Economic Behavior & Organization* 89 (2013), 122–135.
- FRÉCHETTE, G. R., A. SCHOTTER AND I. TREVINO, “Personality, information acquisition, and choice under uncertainty: An experimental study,” *Economic Inquiry* 55 (2017), 1468–1488.
- GREINER, B., “Subject pool recruitment procedures: organizing experiments with ORSEE,” *Journal of the Economic Science Association* 1 (2015), 114–125.
- GRONAU, R., “Wage Comparisons—A Selectivity Bias,” *Journal of Political Economy* 82 (1974), 1119–1143.
- HAMILTON, B. H., N. W. PAPAGEORGE AND N. PANDE, “The Right Stuff? Personality and Entrepreneurship,” *Personality and Entrepreneurship (May 19, 2014)* (2014).
- HECKMAN, J., “Shadow prices, market wages, and labor supply,” *Econometrica: journal of the econometric society* (1974), 679–694.
- HECKMAN, J. J. AND B. E. HONORE, “The empirical content of the Roy model,” *Econometrica: Journal of the Econometric Society* (1990), 1121–1149.
- HERZ, H., D. SCHUNK AND C. ZEHNDER, “How do judgmental overconfidence and overoptimism shape innovative activity?,” *Games and Economic Behavior* 83 (2014), 1–23.
- HEY, J. D. AND C. ORME, “Investigating generalizations of expected utility theory using experimental data,” *Econometrica: Journal of the Econometric Society* (1994), 1291–1326.
- HIRSCH, A. V., “Experimentation and persuasion in political organizations,” *American Political Science Review* 110 (2016), 68–84.
- JOHNSON, J. A., “Measuring thirty facets of the Five Factor Model with a 120-item public domain inventory: Development of the IPIP-NEO-120,” *Journal of Research in Personality* 51 (2014), 78–89.
- JOVANOVIC, B., “Job matching and the theory of turnover,” *Journal of political economy* 87 (1979), 972–990.
- KNIGHT, F., “Risk, uncertainty and profit,” (1921).
- LUNDBERG, S., “The college type: Personality and educational inequality,” *Journal of Labor Economics* 31 (2013), 421–441.

- MANSO, G., "Motivating innovation," *The Journal of Finance* 66 (2011), 1823–1860.
- MARCH, J. G., "Exploration and exploitation in organizational learning," *Organization science* 2 (1991), 71–87.
- MCCLELLAND, D. C., "N achievement and entrepreneurship: A longitudinal study.," *Journal of personality and Social Psychology* 1 (1965), 389.
- MELOSO, D., J. COPIC AND P. BOSSAERTS, "Promoting intellectual discovery: patents versus markets," *Science* 323 (2009), 1335–1339.
- MILLER, R. A., "Job matching and occupational choice," *Journal of Political economy* 92 (1984), 1086–1120.
- RAVEN, J. C. AND J. H. COURT, *Raven's progressive matrices and vocabulary scales* (Oxford Psychologists Press, 1998).
- ROBBINS, H., "Some aspects of the sequential design of experiments," *Bulletin of the American Mathematical Society* 58 (1952), 527–535.
- ROTTER, J. B., "External control and internal control," *Psychology today* 5 (1971), 37–42.
- ROY, A. D., "Some thoughts on the distribution of earnings," *Oxford economic papers* 3 (1951), 135–146.
- SMITH, W. K. AND M. L. TUSHMAN, "Managing strategic contradictions: A top management model for managing innovation streams," *Organization science* 16 (2005), 522–536.
- TUSHMAN, M. L., O. REILLY AND A. CHARLES III, "Organizations: MANAGING EVOLUTIONARY," *California Management Review* 38 (1996), 4.
- WEITZMAN, M. L., "Optimal search for the best alternative," *Econometrica: Journal of the Econometric Society* (1979), 641–654.
- WILLIS, R. J. AND S. ROSEN, "Education and Self-Selection," *The Journal of Political Economy* (1979), S7–S36.