

University of Toronto
Department of Economics



Working Paper 677

Identifying Rule-Based Rationality

By Yoram Halevy and Guy Mayraz

October 25, 2020

Identifying Rule-Based Rationality ^{*}

Yoram Halevy[†] Guy Mayraz[‡]

October 19, 2020

Abstract

The revealed preference methodology allows an observer to infer preferences from choices. This paper extends this fundamental idea by experimentally identifying the preference for basing choices on simple decision rules. Subjects not only make case-by-case portfolio allocations but also design a simple investment rule for selecting portfolios. They then choose between these two decision modes for an additional set of problems. The majority opt for the rule interface, and in most cases, choose a simple investment rule that cannot be rationalized by any simple utility function or accounted for by reductions in decision time or cognitive costs.

^{*}First presented at the D-TEA (Decision: Theory, Experiments and Applications) workshop in June 2016. First version circulated June 2018 under the title “Modes of rationality: act versus rule-based decisions”. Acknowledgments to be added. This study was approved by the University of Melbourne’s Human Research Ethics Committee 1544278.1.

[†]University of Toronto, yoram.halevy@utoronto.ca. Halevy gratefully acknowledges funding from SSHRC Insight Grants 435-2012-0577 and 435-2017-0148.

[‡]University of Sydney, guy.mayraz@sydney.edu.au. Mayraz gratefully acknowledges funding from the Canadian Institute for Advanced Research.

1 Introduction

The dictionary definition of ‘rational’ is “based on or in accordance with reason or logic”.¹ Following the revealed preference tradition (Samuelson, 1938), economists ignore the question of whether choices are based on reason, and focus entirely on whether they accord with logic. Rationality is equated with consistency.

Consistent choices can, of course, be represented by a utility function (Afriat, 1967). Since utility-maximization is a particular rule for making decisions, consistent decision-makers are not only “as-if” utility maximizers (Friedman, 1953), but also “as-if” implementers of a decision rule. Not all decision rules result in consistent choices, but if we restrict attention to those that do, the relationship is one of equivalence. There may thus appear to be no value in studying rule-based choice—after all, the resulting choices can be equally well-described by a utility-function. Yet, nothing in the revealed preference approach guarantees that the recovered utility-function would be a good guide to the decision-maker’s choices in other problems—however closely related.

Whether generalization should or should not be expected would seem to depend on how decision-makers *actually* choose. If the decision-maker has difficulty making choices in some choice setting, she may adopt a decision-rule to guide her choices. As long as the resulting choices are internally consistent, they can be represented by some utility-function. But in the absence of a deep and stable utility-function, we have no grounds to expect that the “as-if” utility-function recovered from observed choices would generalize to other choice-settings

Consider the problem of allocating a budget between several equiprobable Arrow securities. If the price difference between these securities is small, the decision-maker may use a simple decision rule: *purchase the same quantity of all securities*. If the price range is increased, the decision-maker may switch to another simple rule: *purchase only the cheapest security*. Alternatively, the decision-maker may opt for a threshold rule: *purchase the same quantity of*

¹The Pocket Oxford Dictionary of Current English (Oxford University Press, 1984).

those securities that are not more than (say) twice as expensive as the cheapest one. And if, finally, we switch the setting to one in which the budget is allocated between a safe asset and a higher-yielding risky one, the decision-maker may well go back to the original decision rule: *purchase the same quantity of both.*

If we assume that the decision-maker is a utility maximizer, we would conclude from the first set of choices that she is infinitely risk-averse (maxmin preferences). Going by the second choice-set, we would conclude that she is risk-neutral (assuming she opts for the simple *only the cheapest security* rule). And if she opts for the more complex threshold rule, there is no guarantee that we could find *any* utility function to organize her choices,² let alone a simple utility function of the kind empirical economists are likely to consider (Halevy et al., 2018).

If decision-makers really do use decision rules to make their choices, it would be of great value to figure out how they choose their decision rules, and in particular how features of the choice setting affect this choice. In this paper, we take the very first step in this research program: determine experimentally whether decision-makers employ an “as-if” utility function to make their choices, or whether—in the absence of such a utility-function—they select a decision rule that they then follow. We do so in a choice domain that economists are both familiar with and care about: portfolio choice over Arrow securities with equiprobable states.

In their study of a similar environment, Choi et al. (2006, unpublished) find that the portfolio choices of many subjects are most naturally explained as the outcome of simple investment rules. But, of course, any set of internally consistent portfolio choices can also be the outcome of an “as-if” utility function. Our innovation is to offer subjects a choice of interfaces for making their portfolio allocations. In addition to making case-by-case portfolio allocations (as has been done in previous studies), subjects *design* an investment rule for allocating the budget in multiple problems. We identify their choice

²See Footnote 11 for an example of a threshold rule that cannot be represented by any utility function.




Figure 1: In Part 1 of the experiment, subjects experience (in random order) two methods for allocating their budget in portfolio choice problems: making case-by-case decisions and constructing an investment rule. In Part 2, they choose one of these methods for a third set of choice problems.

mode by asking them to decide—for a separate set of portfolio allocation problems—between using an investment rule and allocating their budget directly in each separate problem (Figure 1). The interface for designing investment rules has few degrees of freedom, spanning only a small subset of possible portfolios. Consequently, subjects who have an “as-if” utility function would almost certainly benefit from allocating their budget directly. Nevertheless, a large proportion of subjects—two-thirds of the overall sample—opt to use the investment rule interface. This finding is robust to various design manipulations, and is not explained by a desire to save time or cognitive effort.³

This paper makes three contributions. First, it introduces a novel and intuitive interface for allocating the budget in portfolio choice problems with arbitrarily many Arrow securities. Second, to the best of our knowledge, it is the first study in which subjects design their own rule for making investments. This gives researchers a powerful new tool to study the type of portfolio allocation rules that subjects prefer.⁴ Third, and most importantly, subjects

³We do find, however, significant differences between parts of the sample. We discuss these differences in Section 3.

⁴Axelrod (1980; 1984) conducted tournaments to which leading game theorists and programmers submitted a strategy for the iterative prisoner’s dilemma, thereby obtaining data on preferred procedures for that problem. In a recent study Dal Bó and Fréchet (2019) asked subjects to design a strategy that will best approximate their play in the infinitely repeated prisoners’ dilemma. Although intuitively related, since the strategy method asks subjects to choose an action contingent on various possible histories, it is closer to the case-by-case choices – in which the subject chooses an allocation for different price vector – than

choose between making standard case-by-case allocations and designing an investment rule for allocating their budget. This choice reveals whether subjects have an “as-if” utility function.⁵

Several objections may be raised at this point. One is that subjects who chose the rule interface may have had an “as-if” utility function and would have benefited from using the case-by-case interface, but found it too difficult to use. In order to test this possibility, we asked subjects in the post-experiment survey whether they were able to use the two interfaces to allocate their budget as they wanted. We found that subjects who chose the rule-interface in Part 2 of the experiment rated their ability to use both interfaces significantly higher than the subjects who chose the case-by-case interface.⁶ This observation runs counter to the idea that the choice of the rule-interface was motivated by difficulty using the case-by-case interface. Instead, it supports the opposite interpretation: that at least in some cases, the choice of the case-by-case interface was motivated by the comparative difficulty of using the rule-interface.

A related objection is that subjects may have chosen the rule-interface because the case-by-case interface was too demanding in time or cognitive effort. If this were true, we would expect the choice of the rule-interface to increase with the number of securities. The relative complexity of the case-by-case task increases significantly with this number—measured both by the time required to complete the task and by the proportion of GARP violations. We nevertheless find that the proportion choosing the rule interface is independent of the number of securities. This evidence suggests that, at least in the current environment, complexity (Oprea, forthcoming) is not the driving force for rule-based decisions. As a further test, we replaced the option of reusing the rule from Part 1 of the experiment with the less convenient option of designing a new rule from scratch, and decreased the purported benefit of case-by-case

to the rule choice studied in the current paper.

⁵Recently, [Nielsen and Rehbeck \(2020\)](#) use exogenously assigned decision rules that correspond to some of the classical axioms of expected utility to study whether choices that are inconsistent with expected utility theory are deliberate.

⁶On a scale of 0-100, the mean ratings were 89 and 81 for the case-by-case interface, and 81 and 73 for the rule-interface.

choices by reducing the financial stakes (in part of the sample). Neither manipulation had an impact on the proportion of subjects opting for the investment rule. We also found no relationship between the choice of the investment rule and time spent on the optional survey, suggesting that the subjects who chose the investment rule were in no more of a hurry than the subjects who chose the case-by-case interface.

Finally, some subjects may have had an “as-if” utility function, but were able to find an investment rule that perfectly captures their preferences over the resulting portfolios. We can obtain a bound for the size of this group by examining the rules chosen by the subjects who opt for the rule interface. Only two of the rules that the interface supports have a simple utility function equivalent: allocating the entire budget to the least costly security and allocating the same proportion of the budget to all securities.⁷ We find that 28% of the subjects choose one of these two rules. The rule chosen by the remaining 72% has no simple utility function equivalent. Of course, it is far from clear that subjects with an “as-if” utility function are able to identify an equivalent investment rule—even if such a rule exists.

An entirely different form of evidence is obtained by comparing choices between the two parts of the experiment. The tendency to choose the rule in Part 2 is unrelated to making consistent choices in Part 1. However, when the rule chosen in Part 1 approximates the ranking induced by the case-by-case choices, subjects are more likely to choose the rule interface in Part 2. Moreover, subjects who chose the case-by-case interface in Part 2 exhibited less consistency with their choices in Part 1 than subjects who chose the investment rule interface. Taken together, these findings suggest that the choice of the investment rule interface in Part 2 is deliberate, and that it captures the choice process that subjects actually employ.

Conversely, it is far from clear that those subjects who chose the case-by-case interface have an “as-if” utility function. Some may have had a preference for a rule that they could not design using our restrictive interface. Others may have had trouble using the rule interface effectively. In the optional survey, we

⁷The utility function equivalents are expected value and log utility, respectively.

asked subjects how well they understood the rule interface and how well they were able to use it. Low answers on these two questions are strongly correlated with choosing the case-by-case interface in Part 2 (almost all subjects found the case-by-case interface straightforward). The proportion of subjects who chose the rule interface in Part 2 should probably be seen as a lower bound on the proportion who prefer making rule-based decisions. Finally, the choice of the case-by-case interface does not prove that a decision-maker has an “as-if” utility function. Some decision-makers may prefer to retain direct control over their portfolio allocations, either because they are unsure which investment rule they prefer, or because of a non-instrumental preference for control.

The notion of procedural rationality was introduced into economics by Herbert Simon ([Simon, 1955](#)). Simon’s focus was on situations in which decision-makers have a clear and well-defined goal that can be modeled using a utility function, but where finding the optimal solution is difficult. As an example, a manager choosing the price for a new product is presumably looking to maximize profits. But since this is hard, the manager may use a cost-plus rule-of-thumb, rather than attempt to estimate the demand curve and solve for the profit maximizing price.⁸ Kahneman and Tversky’s heuristics and biases program ([Tversky and Kahneman, 1974](#); [Kahneman et al., 1982](#)) is likewise focused on problems in which people face a difficult task, and handle it by using some heuristic procedure.⁹ These heuristics lead to the behavioral biases the analyst observes. Gigerenzer is often seen as an opponent of Kahneman and Tversky, but in this respect his notion of heuristics is similar.¹⁰ Aumann’s ([2008](#); [2019](#)) “rule-rationality” describes behavior that is optimal in most (regular) cases, but may be sub-optimal in contrived scenarios. Aumann suggests that evolutionary forces have shaped these rules. [Halevy and Feltkamp \(2005\)](#) show how rule-rationality can account for ambiguity aversion, and [Heller and Winter \(2016\)](#) apply it to game-theoretic environments.

The procedures we consider in this paper are different. They are not the

⁸This observation about cost-plus pricing was first made by [Hall and Hitch \(1939\)](#).

⁹For example, when estimating how frequent an event is, decision-makers may use the heuristic that common events are usually easier to bring to mind.

¹⁰For an overview of Gigerenzer’s approach to heuristics see [Gigerenzer et al. \(1999\)](#).

result of the difficulty in finding a solution to a well-defined problem, but of the absence of one. In the absence of a fully-specified utility-function over portfolios or consumption streams, rationality calls for adopting some reasonable decision rules and basing choices on them. Two well-known rules are arguably examples of such rule-based rationality. [Benartzi and Thaler \(2001\)](#) document the use of the $1/n$ heuristic in defined contribution saving plans. Rubinstein suggests that people use a similarity procedure in binary choices between simple lotteries and temporal allocation problems, and argues that this procedure accounts for Allais-type and temporal-biased choices ([Rubinstein, 1988, 2003](#)). Both procedures can be seen as a way of making a rational decision in the absence of an “as-if” utility function.

2 Experimental Design

The setting is of a classical portfolio allocation problem. Subjects divide a budget among several Arrow securities. Each security pays in a single state of the world, and all states are equally likely. The securities differ in price, creating a trade-off between risk and return. Expected payoff is maximized by investing the entire budget in the least expensive security, while the worst-case payoff is maximized by purchasing the same quantity of each security.

Rather than explain Arrow securities to subjects, we describe them as “companies” that compete against each other, and have the same chance of “winning.” Subjects use their budget to purchase shares in these companies and are paid in proportion to the number of shares in the company that wins. Subjects are presented with a total of 30 problems. The budget is the same in all problems (denominated in Experimental Currency Units, ECUs), but the price of shares varies both between companies and across problems. The shares of Company #1 are always the least expensive, and their price ranges from 1 to 2 ECUs/share. The shares of the most expensive company cost up to 4 ECUs/share. One of the 30 problems is randomly selected for payment after the experiment concludes. The budget for each problem has to be spent in its entirety on the companies in that problem.

The number of companies in each problem can be two, three, or six. This number varies across conditions but is held constant within conditions. Each subject is therefore presented with 30 problems that differ only in share prices. The 30 problems are divided into three sets: 1-10, 11-20, and 21-30. In Part 1 of the experiment, subjects try out two different methods for allocating the budget: (i) *case-by-case*, and (ii) *investment rule*. Half the subjects use the case-by-case method for problems 1-10 and the investment rule for problems 11-20; the other half experience these two methods in the opposite order. In Part 2 of the experiment, subjects choose one of these two methods for allocating the budget in problems 21-30 (Figure 1). Our focus is on this choice, as it identifies the decision maker’s choice mode.

The experiment starts with an explanation of the investment problems and payment scheme, followed by a quiz that tests subjects’ understanding of these instructions. Both the case-by-case and investment rule tasks in Part 1 are preceded by a demanding tutorial in which subjects learn to use the relevant interface and go through a sequence of exercises that test their understanding of its capabilities. After allocating the budget in all 30 problems, subjects are presented with an optional survey that asks for demographic information and for their understanding of the different investment methods. Finally, subjects are informed of the problem and company selected for payment, of their choice in that problem (whether directly via the case-by-case interface or indirectly via the investment rule), and of the resulting payment. For further information, refer to Appendix A, which includes screenshots of the different parts of the experiment.

We track the content and timing of all subject actions, including intermediate choices. For example, whenever a subject drags a slider to alter the case-by-case allocation in a particular problem, we have a record of the resulting allocation and the time in which it was made. Most importantly, we have a record of the time subjects spent on each section of the experiment, including the tutorials and the survey.

2.1 Case-by-case interface

Subjects use the case-by-case interface to allocate the budget directly in each problem. Since we wanted the choice of the rule interface in Part 2 to be meaningful, it was important to make the case-by-case interface consistent across problems that differ by the number of companies, and as easy to use as possible. Previous work (Choi et al., 2007; Halevy et al., 2018) used a “budget line” interface that works well for allocating the budget between two securities but is challenging to use with three securities, and cannot be scaled to problems with four or more securities. We thus implemented an intuitive new interface that allowed us to present problems with two, three, or six securities in a consistent way.

Subjects in a condition with n companies allocate their budget (denominated in ECUs) by moving n sliders: one for each company (Figure 2). The sliders are ordered by share price from the least expensive to the most expensive company. Subjects purchase shares in a company by moving its slider to the right and reduce their holding by moving the slider to the left. The range is limited by the budget so that subjects can allocate the budget in its entirety by simply pushing one of the sliders to the right as far as it would go.

Based on previous research, we conjectured that some subjects would like to purchase the same number of shares in two or more companies. In order to make this easy, whenever subjects push the slider of a company c to the right, the interface automatically also purchases shares in any company c' with a lower share price, thereby ensuring that the number of shares in c and c' remains the same. Similarly, whenever subjects push the slider of c' to the left, the interface automatically also sells shares in c . Subjects can thus purchase a baseline of x shares in the first n companies by pushing the n 'th slider to x . Purchasing the same number of shares in all companies is as easy as pushing the last slider until the budget is exhausted.

This arrangement limits the subjects' freedom by making it impossible for them to purchase more shares in a company with a higher share price. Since all companies have the same chance to “win”, such a portfolio would be first-order stochastically dominated. We thus lose the opportunity of observing

Investment problems 1-10

Your task is to make case-by-case investment decisions in problems 1-10. Use the ◀ and ▶ arrows to move between problems. Submit your choices once you are happy with your budget allocations in all the problems.

The screenshot shows the interface for 'Investment problem #1'. It features four horizontal sliders with labels on the left and values on the right. The first slider is for a company with a price of 1.10, showing 155 shares. The second slider is for a company with a price of 1.30, also showing 155 shares. The third slider is for a company with a price of 3.00, showing 0 shares. The fourth slider is for the 'Budget', showing 228 ECUs. Below the sliders are two buttons: 'Review instructions' and 'Submit your choices'.

Company	Price	Shares
Company 1	1.10	155
Company 2	1.30	155
Company 3	3.00	0
Budget	-	228 ECUs

Figure 2: The case-by-case allocation screen in the 3 companies condition.

such transparent violations of first-order stochastic dominance, but gain a much more user-friendly case-by-case interface that makes the choice of the rule interface in Part 2 considerably more significant. We can still, of course, observe GARP violations, which occur when the choices subjects make in different problems cannot be rationalized by any single preference relation.

2.2 Investment rule interface

Subjects use the investment rule interface to select a rule for allocating the budget in a generic portfolio choice problem. This rule is then implemented in all the problems in the set. When designing the rule, subjects can see its impact in each of the 10 problems in the set. This helps subjects understand how the rule works concretely in the problems in which it will be implemented. Subjects cannot, however, choose different rules for different problems. To the best of our knowledge, this is the first time that subjects are presented with such a task, in either a lab or a field setting.

The investment rule interface (Figure 3) includes the following degrees of freedom: (i) the number of companies the subject invests in, (ii) how to invest the budget in these companies, and (iii) an optional baseline number of shares in all companies, which guarantees the subject a sure payment. Starting with the first degree of freedom, subjects can either choose to invest in the same

number of companies in each problem, or they can let this number vary as a function of the relative share price. Subjects do this by specifying a threshold price level relative to that of the company with the least expensive share price. This selects only the companies at or below the threshold – a number which may be small in some problems (perhaps only Company #1) and large in others (perhaps all companies). The second degree of freedom determines how the budget is allocated among the companies selected in the first part. Subjects have a binary choice between purchasing the same number of shares in these companies and allocating the same proportion of the budget to them (that is, the same number of ECUs). This choice affects the resulting allocation whenever the budget is allocated among two or more companies. The final degree of freedom is an optional baseline number of shares purchased in all the companies. If this option is selected, the budget is used first for purchasing this baseline, and the first two parts of the rule are applied only to the residual budget. Figure 3 illustrates such a baseline rule.

In designing the investment rule interface, we relied both on previous experimental work and on our intuition. [Choi et al. \(2006\)](#) study case-by-case choices and classify subjects into archetypes based on their chosen portfolios. These types inspired our two first degrees of freedom. [Halevy et al. \(2018\)](#) note the importance of the baseline for some subjects (and also its implied non-convexity if the researcher uses a simple parametric form of preferences). This observation led us to include a baseline as a third degree of freedom. The resulting interface by no means includes all plausible investment rules. As an example of a plausible rule that cannot be implemented using our interface, consider using a given percentage of the budget to guarantee a safe return, and investing the rest in Company #1. There are, of course, many other possibilities.

Several of the investment rules that can be implemented using our interface induce choices that can be represented by simple utility functions: investing in the cheapest company is equivalent to risk-neutrality; purchasing an equal number of shares in all companies is equivalent to infinite risk aversion; investing an equal number of ECUs in all the companies is equivalent to maximizing

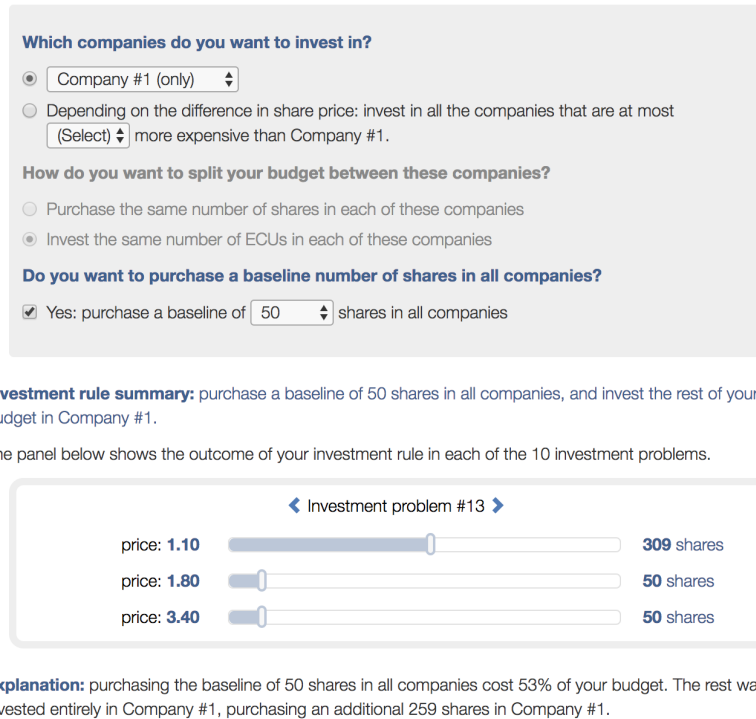


Figure 3: The investment rule interface in the 3 companies condition.

expected log utility. Other investment rules have no simple utility function equivalent. A few investment rules can even induce portfolios that are inconsistent with GARP, and therefore cannot be rationalized by any utility function.¹¹

2.3 The second part of the experiment

After experiencing both methods of allocating the budget in Part 1, subjects are asked in Part 2 which of the two investment methods they would like to

¹¹As an example, consider a rule in the 3 companies condition that allocates equal shares to companies priced at up to twice the price of Company #1, and let the budget be 450. With a price vector $P_1 = (6, 12, 14)$, the allocation is $X_1 = (25, 25, 0)$. With a price vector $P_2 = (8, 9, 16)$, the allocation is $X_2 = (450/33, 450/33, 450/33)$. The expenditure of X_1 under P_2 is $425 < 450$; hence, X_2 is strictly revealed preferred to X_1 . The expenditure of X_2 under P_1 is $436.36 < 450$; hence, X_1 is strictly revealed preferred to X_2 . GARP is violated and Afriat's CCEI is less than 1.

use for problems 21-30. Subjects who choose the case-by-case method use the interface described in Section 2.1. Subjects who choose the investment rule method use the interface described in Section 2.2.

2.4 Optional survey

After completing the task, subjects were presented with an optional survey. In addition to basic demographic details, the survey includes questions on the subjects' understanding of the case-by-case and rule interfaces and their ability to use them to allocate the budget effectively. The questions on the rule interface were (i) 'Did you have a good understanding of the investment rule interface?', and (ii) 'Were you able to use the interface to allocate the budget the way you wanted to?'. Both questions were answered using a slider ranging from 'No' to 'Yes' with answers coded as a real number between 0 and 1. Other (mainly open) questions elicited the reasons for the choices subjects made. For further details, see Appendix A.

By making the survey optional, we created an instrument for identifying those subjects who were particularly keen to save time. Such subjects could be expected to skip some or all questions, and to finish the survey faster than other subjects. We use this to investigate the link between the choice in Part 2 and the desire to save time.

2.5 Conditions and experimental details

The experiment was conducted between January and April 2016. Subjects ($N = 324$) were recruited either from Amazon Mechanical Turk workplace (mTurk, $N = 198$) or from the University of Melbourne's experimental subject pool (lab, $N = 126$).¹² Only US-based mTurk workers with a minimum of 1000

¹²One additional lab subject and 20 additional mTurk subjects signed up for the experiment and clicked through the instructions, but did not finish. The lab subject dropped out during the rule tutorial. The 20 mTurk drop-outs occurred during the rule tutorial (15), case-by-case tutorial (2), case-by-case task (2), and survey (1).

completed tasks and a minimum approval rate of 97% were allowed.¹³ mTurk workers received a variable payment of up to 12 US dollars via the mTurk bonus mechanism, as well as a \$3 fee for completing the experiment. The average total payment was \$4.13. Lab subjects received either a show-up fee of A\$7 and a variable payment of up to A\$90 ($N = 75$), or a show up fee of A\$13 and the same variable payment as mTurk workers ($N = 51$). The average total payment in the two groups of lab subjects was almost identical: A\$14.19 and A\$14.16.¹⁴ In order to minimize differences with mTurk workers, lab subjects were allowed to leave and collect their payment after finishing their own task. We chose to include both mTurk workers and students in order to get a sense of the distribution of the preference for rule-based decisions in two very different subject pools.

Subjects were randomly assigned to conditions with $n \in \{2, 3, 6\}$ companies ($N = 112, 107, \text{ and } 105$), and independently randomly assigned to case-by-case in problems 1–10 and investment rule in problems 11–20 ($N = 166$) or the opposite order ($N = 158$). The budget was proportional to the number of companies so that the expected payoff for risk-neutral choices was approximately the same. There were two versions of the rule option in Part 2. Some subjects ($N = 144$) were asked to choose between case-by-case decisions and the rule they chose earlier in the experiment (after observing the induced allocations); others ($N = 180$) were asked to choose between case-by-case decisions and designing a new rule for the remaining problems. These random assignments were independent of the subject group (mTurk, lab, or lab with mTurk stakes).

The 30 problems in each condition were drawn from the same distribution. For about a sixth of subjects ($N = 53$) the problems were all different, but

¹³Similar restrictions were used in [Berinsky et al. \(2012\)](#), [Freeman et al. \(2019\)](#), and [Freeman and Mayraz \(2019\)](#). Amazon holds this information about mTurk workers, and enables experimenters to limit the sample by these criteria. The mTurk sessions were completed well before the start of the recent (2018) concern with bots posing as mTurk workers. It is also notable that the case-by-case and (especially) rule tutorials were highly demanding and could not possibly be completed by any general purpose bot. Finally, 194 of the 198 mTurk subjects completed open ended survey questions on how they made their decisions, and in all 194 cases the answers were relevant to the question.

¹⁴The exchange rate at the time was roughly $\text{A\$1} = \text{US\$0.80}$.

for the remaining subjects ($N = 271$) we repeated the same 10 problems in problems 1-10, 11-20, and 21-30, but in a different order in each set. Repeating the problems allowed us to make a direct comparison of the allocations subjects made in different parts of the experiment. Subjects were not told that problems repeat, and the scrambling of problems between sets would have made it very difficult for subjects to notice the repetition. Subjects were also not aware of the variations between treatments. The printed forms and computer interface of each subject included only the details for the treatment that the subject was assigned to.

3 Results

Our major finding (Table 1) is that two-thirds of all subjects chose to make their budget allocations using an investment rule. About the same proportion of rule-choosers was observed in conditions with two, three, and six companies. This is despite the substantial decrease in GARP scores—evidence of the increased difficulty of choosing portfolios consistently in conditions with many companies (a measure of “complexity”, see [Oprea, forthcoming](#)).

Table 1: Key results by the complexity of the portfolio allocation task.

Complexity of condition	Subjects in condition	Consistent with GARP	Chose rule in Part 2
2 companies	112	78%	68%
3 companies	107	58%	64%
6 companies	105	50%	70%
All conditions	324	62%	67%

The factors associated with this choice are explored in Table 2. Columns report the coefficients and standard errors in probit regressions, where the dependent variable is whether the subject chooses to use the investment rule interface to allocate the budget in Part 2 of the experiment. In Column (1) we see again that this choice is not associated with the number of companies.

The order in which subjects experienced the case-by-case and investment rule interfaces also appears irrelevant, as does the difference in stakes between the mTurk and lab samples (recall that we gave the mTurk incentives to a subset of the lab subjects). However, controlling for the difference in stakes, mTurk subjects were significantly more likely than lab subjects to choose the investment rule interface in Part 2 (79% vs. 49%). Among the subjects who chose the rule interface in Part 2, there was a small difference in the proportion that can be rationalized by a simple utility function:¹⁵ 31% of mTurkers (30% unconditionally) and 21% among lab subjects (25% unconditionally). In other words, 54% of mTurk subjects and 39% of lab subjects selected a rule that cannot be represented by a simple utility function in Part 1 and then opted for the rule interface in Part 2.

We do not know why mTurk subjects were more likely to choose the investment rule interface. Table 3 reports some demographic data on the mTurk and student samples. mTurk subjects were older and more likely to be male, but neither factor is significantly correlated with choosing the rule in Part 2. Our mTurk subjects were based in the US, whereas our lab subjects were based in Australia. There were also differences in education, income, and other unobserved factors. Our best conjecture is that the rule interface is not trivial to master, and that mTurkers were, on average, more confident in using it. One possible clue is that mTurk subjects were considerably faster at completing the rule tutorial, spending a median of 4.7 minutes on it vs. 6.2 minutes for students. By comparison, neither group showed much difficulty with the case-by-case tutorial, spending on it a median of only 2.0 and 2.1 minutes. We also know that (self-reported) confidence in using the rule interface was strongly correlated with choosing the rule in Part 2 (Section 3.4). Hence, it is possible that the preference for rule-based choices was about as high in both groups, but that some lab subjects chose the case-by-case interface because they were not confident with the rule interface. As for why mTurkers were

¹⁵As we note in the Introduction, this category includes the rule that allocates the entire budget to Company #1 (equivalent to maximizing expected value or any risk seeking utility function) and the rule that allocates the same proportion of the budget to all companies (equivalent to maximizing log utility).

Table 2: Factors associated with the choice of investment rule in Part 2.

	(1)	(2)	(3)
Number of companies in condition	0.03 (0.04)	0.02 (0.05)	0.00 (0.05)
Case-by-case in problems 1-10	0.04 (0.15)	0.07 (0.16)	0.14 (0.18)
Redesign rule in part 2	-0.23 (0.17)	-0.22 (0.17)	-0.01 (0.19)
Amazon Mechanical Turk stakes	0.14 (0.25)	0.10 (0.26)	0.03 (0.29)
Amazon Mechanical Turk sample	0.72*** (0.22)	0.70*** (0.22)	0.72** (0.29)
Part 1 case-by-case choices violate GARP		0.05 (0.16)	0.24 (0.18)
Drop in Afriat CCEI when adding rule allocations		-3.94 * (2.10)	-4.99 ** (2.30)
Rule: company #1 only (with or w/o baseline)		0.62** (0.24)	0.69*** (0.26)
Rule: equal number of ECUs (rather than shares)		0.43** (0.19)	0.53** (0.21)
Rule includes a baseline		-0.18 (0.17)	-0.25 (0.19)
Age			-0.02 * (0.01)
Male			0.02 (0.17)
Can use case-by-case interface effectively (self-reported)			-0.27 (0.67)
Can use rule interface effectively (self reported)			1.91*** (0.61)
Time in Part 1 case-by-case allocations (minutes)			-0.04 (0.04)
Time in optional survey (minutes)			-0.05 (0.05)
Observations	324	324	301
Pseudo R^2	0.080	0.105	0.179

Separate probit regressions in each column. Dependent variable: whether the subject chose the investment rule in Part 2. Standard errors in parentheses. Statistical significance indicators: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

better at mastering the rule interface, we note that mastering a great variety of online tasks is an integral part of mTurk work. It would not be surprising if this experience made it easier for mTurk workers to master the unfamiliar rule interface in our experiment.¹⁶

Table 3: Differences between mTurk and student samples.

	mTurk	Students	Combined
Number of subjects	198	126	324
Age (median)	33	21	32
Male (%)	59	47	55
Chose investment rule (%)	79	49	67

3.1 Consistency of choices in Part 1

We now explore whether there is a relationship between the choice subjects make in Part 2 and the consistency of their case-by-case choices in Part 1. The usual yardstick for the coherence of case-by-case choices is consistency with the Generalized Axiom of Revealed Preference (GARP, [Afriat, 1967](#)). If a subject fails this test, several inconsistency indices can be used to measure the severity of the inconsistency. We opted to use Afriat’s Critical Cost Efficiency Index (CCEI, [Afriat 1972, 1973](#)). Remember that if choices satisfy GARP, there exists a utility function that rationalizes these choices. [Table 1](#) reports the number of subjects who pass the GARP test as a function of the number of states (companies). As the number of states increases, the proportion of subjects whose choices are consistent with GARP drops from 78% (2 companies) to 50% (6 companies). This confirms the intuition that complexity increases with the number of states ([Oprea, forthcoming](#)). Despite the increase in complexity, it appears from both [Table 1](#) and from the probit regressions in [Table 2](#) that there is no relationship between the number of companies and the propensity to use an investment rule in Part 2.

¹⁶In addition to the experience effect, there may also be a selection in mTurk for those individuals who are particularly good at quickly mastering unfamiliar computer interfaces.

We further explored the consistency between the portfolios chosen in Part 1: the portfolios induced by the investment rule and the portfolios chosen directly using the case-by-case interface (recall that subjects experienced both methods of portfolio allocation in Part 1). Afriat’s CCEI measures how far (if at all) are the case-by-case portfolios from being rationalized by a consistent preference ranking. When we add the portfolios induced by the rule, Afriat’s CCEI does not change if the rule is consistent with the ranking (or approximate ranking, see [Halevy et al., 2018](#)) induced by the case-by-case portfolios. We would, however, expect a substantial drop in Afriat’s CCEI if the ranking implied by the rule is very different from the one that rationalizes the case-by-case choices. As can be seen in Table 2, the higher the drop in Afriat’s CCEI is, the less likely is the subject to choose the rule interface in Part 2. In other words, the more consistent the case-by-case allocations and the allocations induced by the rule designed in Part 1, the more likely was the subject to choose the rule in Part 2.

3.2 The investment rules that subjects choose

Table 4 reports the distribution of investment rules that subjects choose in Part 1 of the experiment. We classify subjects (i) by the number of companies they invest in: Company #1 only, some companies (the number can be a function of how expensive are the shares of other companies relative to those of Company #1), or all companies; (ii) how the budget is allocated among these companies (an equal number of shares or an equal number of ECUs); and (iii) whether they purchase a baseline number of shares in all companies to guarantee themselves a baseline income.

Almost 24% of subjects chose to invest in the cheapest company (41.6% of them also decided to purchase a baseline of shares in all companies). 51% of subjects chose to allocate an equal number of ECUs to some or all companies, and the remaining 25% chose to allocate an equal number of shares to some or all companies. All in all, 51% of subjects chose an investment rule that guarantees a known positive income in all states and whatever the prices may

Table 4: The investment rules subjects choose in Part 1.

Number of companies	Equalize ECUs or shares ^a	% subjects using rule	% subjects adding a baseline to rule ^b
#1 only	N/A	23.8	41.6
Some	ECUs	30.0	45.4
Some	Shares	10.5	32.4
All	ECUs	21.0	45.6
All	Shares	14.8	N/A ^c

^a When allocating the budget to more than one company, subjects can choose between equalizing the number of ECUs invested in each of these companies or equalizing the number of shares they buy.

^b This column indicates the percentage of subjects using the rule in that row who chose to add a baseline to the rule.

^c This rule does not change if a baseline is added.

be—whether by choosing a rule with a baseline or by fully insuring themselves.

The rule chosen in Part 1 was significantly associated with the Part 2 choice. Table 2 reveals that both subjects who chose in Part 1 a rule that invested only in the cheapest company (with or without a baseline) and subjects who chose in Part 1 a rule in which the number of shares bought is inversely related to their price (by equalizing ECUs) were significantly more likely to choose the investment rule interface in Part 2.

3.3 Consistency of choices between Parts 1 and 2

We examined the consistency of choices between the two parts of the experiment in the 180 subjects who were asked in Part 2 to choose between case-by-case allocations and designing a *new* decision rule. The test involved choices in 20 problems: the 10 case-by-case allocations that the subjects chose in Part 1 of the experiment, and the 10 allocations they made in Part 2 (whether directly using the case-by-case interface or indirectly through the choice of a new decision rule). 114 of these 180 subjects chose the rule (63%).

Subjects with severe GARP violations in these 20 problems—an Afriat CCEI of 0.85 or below—were highly unlikely to choose the investment rule in Part 2 (only 1 out of 9 subjects did so). Subjects who violated GARP, but

had an Afriat CCEI above 0.85, were somewhat more likely than average to choose the investment rule (65 out of 91, or 71%), and those who were fully consistent with GARP were somewhat less likely to choose the rule (48 out of 80, or 60%). It is also notable that the reduction in Afriat's CCEI as a result of including the choices made in Part 2 was twice as big on average (0.022 vs. 0.011) in the group that chose the case-by-case interface in Part 2 than in the group that chose the rule interface. Large drops in CCEI were associated with choosing the case-by-case interface: 9 of 16 subjects whose Afriat's CCEI dropped by more than 0.05 chose to make their budget allocations in Part 2 using the case-by-case interface—a substantially higher proportion than the 1/3 average across the entire sample .

3.4 Response to the optional survey

Column (3) of Table 2 adds answers from the optional non-incentivized survey. The Part 2 choice is unrelated to gender, but older subjects are perhaps a little less likely to choose the investment rule. Self-reported understanding of the case-by-case interface is insignificant, but the corresponding self-reported understanding of the rule interface is highly significant. In fact, it is the single strongest correlate of choosing the investment rule interface in Part 2.

4 Extrinsic rationales for choosing the rule

We may question whether choosing the investment rule interface in Part 2 reveals a true preference for making choices based on rules. Perhaps subjects choose the rule interface in order to save the time or cognitive effort involved in making case-by-case allocations, which—despite our effort in making the case-by-case interface easy to use—may be higher than that of the rule interface. Other things being equal, such subjects would prefer using the case-by-case interface, but the associated costs are too high. This, indeed, is one way of explaining the lower proportion of rule choices in the lab, where the time and cognitive costs of using the case-by-case interface are the same, but the typical

dollar stakes are much higher.

We have several methods for testing this hypothesis. Starting with the benefit side, we include lab conditions in which the stakes are the same as in mTurk conditions (but with a higher participation fee). In these conditions, the hypothesized benefit of using the case-by-case interface is lower, so we should expect a higher proportion of rule choices. Contrary to this hypothesis, the proportion of rule choices in this group is the same as for other lab subjects (‘Amazon Mechanical Turk stakes’ in Table 2). We also include conditions in which subjects who choose the rule in Part 2 cannot reuse the rule they designed in Part 1, but have to design a new rule from scratch (‘Redesign rule in Part 2’ in Table 2). In these conditions, the advantage in time and cognitive effort of choosing the rule in Part 2 is smaller, but the proportion of rule choices is again unchanged.

Moving on to the cost side, we can exploit the fact that both the time and cognitive complexity of using the case-by-case interface is increasing in the number of companies. The median time for completing the case-by-case task increases from 2 minutes and 15 seconds in two companies conditions to 3 minutes and 5 seconds in six companies conditions. The increased cognitive complexity is evidenced by the drop in the proportion of subjects whose case-by-case portfolio allocations are consistent with GARP (Table 1). The rule interface, by contrast, does not change with the number of companies, and completion times are also unchanged (a median of 1 minute and 5 seconds in two companies conditions and a median of 1 minute and 2 seconds in six companies conditions). Thus, if the time and cognitive costs hypothesis is correct, we should expect the proportion of subjects choosing the rule interface in Part 2 to increase with the number of companies. But as can be seen in both Table 1 and Table 2, there is no such increase. It is also notable that on an individual level, neither the time spent on the case-by-case allocations (‘Time in Part 1 case-by-case allocations’) nor the subject’s own difficulty of making consistent allocations (as measured by consistency with GARP) is correlated with choosing the rule in Part 2.

Finally, we deliberately made the survey at the end of the experiment op-

tional and measured the time subjects choose to spend on it before leaving the experiment. If subjects who choose the rule in Part 2 are motivated by a desire to minimize effort and finish the experiment quickly, we would expect them to rush through the survey and skip questions. Contrary to this prediction, almost all subjects (305 of 324) completed all the main questions in the survey, and the 19 subjects who skipped some of these questions were actually less likely (53%) to choose the rule than subjects who answered all these questions (68%). The time subjects spent on the optional survey was not significantly associated with the choice in Part 2 (‘Time in optional survey’ in Table 2).¹⁷

5 Discussion

Two-thirds of our subjects prefer making their allocations in Part 2 using an investment rule. This choice does not appear to be related to the time and cognitive costs of using the case-by-case interface. Instead, it reveals what seems to be a genuine preference for basing portfolio allocations on a simple investment rule.

As is clear from Table 2, the most important correlate of the Part 2 choice is the subject’s self-reported rating of her understanding of the rule interface and ability to use it to allocate the budget the way she wanted to (‘Can use rule interface effectively’ in Table 2).¹⁸ A plausible interpretation is that virtually all subjects felt confident using the case-by-case interface, but some subjects did not feel comfortable with the rule interface. Such subjects understandably opted to make case-by-case decisions in Part 2. If this interpretation is correct, it follows that the preference for using an investment rule among subjects who

¹⁷mTurk and student subjects spent approximately the same time in the optional survey: 3.1 vs 3.25 minutes on average, respectively.

¹⁸This is the mean of two survey questions: ‘Did you have a good understanding of the investment rule interface?’ and ‘Were you able to use the interface to allocate the budget the way you wanted to?’. Both questions were answered using a slider ranging from ‘No’ to ‘Yes’ with answers coded as a real number between 0 and 1. Both answers are strongly (and strongly statistically significantly) positively associated with the choice of the rule interface in Part 2, as well as with each other. Because of this high correlation, we use their mean in Table 2.

understood the rule interface is higher than the 2/3 figure in the entire sample.

The consistency of case-by-case allocations does not by itself indicate the mode of rationality. Consistent allocations could either be the result of maximizing a complete and transitive ranking over portfolios, or the result of implementing a decision rule. A better indication is offered by the consistency between the decision rule and the portfolios chosen through the case-by-case interface. Such consistency suggests that the subject employs rule-based rationality in her choices.

In conclusion, we believe it is appropriate to recall the limits of the revealed preference approach we use to identify the preference for rule-based choice. A subject who chooses to base her portfolios in Part 2 on an investment rule chooses *as-if* she uses rule-based rationality. Extrinsic motivations cannot be ruled out, though we believe that we were able to exclude the most obvious ones: the desire to save time and cognitive effort. This paper can be viewed as opening the detailed empirical study of the foundations of rule rationality, and the mapping of the set of rules people use to guide their choices. The role of bounded rationality and incomplete preferences seem to us particularly promising on the theoretical side, while understanding the association between documented behaviors and rule-based reasoning is a particularly important avenue of empirical research.

References

- AFRIAT, S. N. (1967): “The construction of utility functions from expenditure data,” *International Economic Review*, 8, 67–77.
- (1972): “Efficiency estimates of production functions,” *International Economic Review*, 13, 568–598.
- (1973): “On a system of inequalities in demand analysis: an extension of the classical method,” *International Economic Review*, 14, 460–472.
- AUMANN, R. J. (2008): “Rule-rationality versus act-rationality,” research re-

- port 497, Center for the Study of Rationality, The Hebrew University of Jerusalem.
- (2019): “A synthesis of behavioral and mainstream economics,” *Nature Human Behavior*, 3, 666–670.
- AXELROD, R. (1980): “Effective choice in the Prisoner’s Dilemma,” *Journal of Conflict Resolution*, 24, 3–25.
- AXELROD, R. M. (1984): *The Evolution of Cooperation*, Basic Books Inc.
- BENARTZI, S. AND R. H. THALER (2001): “Naive diversification strategies in defined contribution saving plans,” *American economic review*, 91, 79–98.
- BERINSKY, A. J., G. A. HUBER, AND G. S. LENZ (2012): “Evaluating online labor markets for experimental research: amazon.com’s mechanical turk,” *Political Analysis*, 20, 351–368.
- CHOI, S., R. FISMAN, D. GALE, AND S. KARIV (2006): “Substantive and procedural rationality in decisions under uncertainty,” .
- (2007): “Consistency and heterogeneity of individual behavior under uncertainty,” *American Economic Review*, 95, 1921–1938.
- DAL BÓ, P. AND G. R. FRÉCHETTE (2019): “Strategy Choice In the Infinitely Repeated Prisoner’s Dilemma,” *American Economic Review*, 109, 3929–52.
- FREEMAN, D., Y. HALEVY, AND T. KNEELAND (2019): “Eliciting risk preferences using choice lists,” *Quantitative Economics*, 10, 217–237.
- FREEMAN, D. AND G. MAYRAZ (2019): “Why choice lists increase risk taking,” *Experimental Economics*, 22, 132–154.
- FRIEDMAN, M. (1953): “The methodology of positive economics,” in *Essays in Positive Economics*, University of Chicago Press.
- GIGERENZER, G., P. M. TODD, AND THE ABC RESEARCH GROUP (1999): *Simple Heuristics That Make Us Smart*, Oxford University Press.

- HALEVY, Y. AND V. FELTKAMP (2005): "A Bayesian Approach to Uncertainty Aversion," *The Review of Economic Studies*, 72, 449–466.
- HALEVY, Y., D. PERSITZ, AND L. ZRILL (2018): "Parametric recoverability of preferences," *Journal of Political Economy*, 126, 1558–1593.
- HALL, R. L. AND C. J. HITCH (1939): "Price theory and business behaviour," *Oxford economic papers*, 12–45.
- HELLER, Y. AND E. WINTER (2016): "Rule rationality," *International Economic Review*, 57, 997–1026.
- KAHNEMAN, D., P. SLOVIC, AND A. TVERSKY, eds. (1982): *Judgment under uncertainty: Heuristics and biases*, Cambridge University Press.
- NIELSEN, K. AND J. REHBECK (2020): "When Choices are mistakes," .
- OPREA, R. (forthcoming): "What makes a rule complex," *American Economic Review*.
- Oxford University Press (1984): *The Pocket Oxford Dictionary of Current English*.
- RUBINSTEIN, A. (1988): "Similarity and decision-making under risk (Is there a utility theory resolution to the Allais paradox?)," *Journal of Economic Theory*, 46, 145–153.
- (2003): "'Economics and Psychology'? The case of Hyperbolic Discounting," *International Economic Review*, 44, 1207–1216.
- SAMUELSON, P. A. (1938): "A note on the pure theory of consumer's behaviour," *Economica*, 5, 61–71.
- SIMON, H. A. (1955): "A behavioral model of rational choice," *The quarterly journal of economics*, 69, 99–118.
- TVERSKY, A. AND D. KAHNEMAN (1974): "Judgment under uncertainty: heuristics and biases," *Science*, 185, 1124–1131.

A Experimental design details

The experiment includes the following sections:

A.1 General instructions

The general instructions (Figure 4) explain the overall structure of the experiment, the nature of the investment problems, and the payment scheme.

A.2 Quiz

The quiz (Figures 5 and 6) tests subjects' understanding of the instructions, and in particular, how their payment is related to the number of shares they have in each of the companies.

A.3 Task

In the task, subjects make a total of 30 portfolio allocations divided into three groups. Depending on a random draw, problems 1-10 are allocated using the case-by-case interface and problems 11-20 using the rule interface, or the order is reversed. In either case, subjects choose which interface to use for problems 21-30. In the paper, the first two groups of problems are referred to as 'Part 1' and the third group (and especially the choice between the case-by-case and rule interfaces) is referred to as 'Part 2'. In the instructions to subjects, however, these three groups were referred to as 'Part 1', 'Part 2', and 'Part 3'. Readers should keep this difference in terminology in mind when they read this section and examine the screenshots.

A.3.1 Budget allocation in problems 1-10

The method used for allocating the budget in these problems was randomly chosen to be either case-by-case or the investment rule. The screenshots included here assume case-by-case decisions in problems 1-10 and an investment rule decisions in problems 11-20.

The budget allocation task was preceded by instructions and a tutorial (Figure 7). The tutorial walks the subject through the use of the interface, using examples selected to illustrate its power. The case-by-case tutorial includes four tasks that have to be successfully completed before the next task is revealed (Figure 7 shows the screen after all four tasks were completed). It is only after the successful completion of these four tasks that the subject can proceed to allocate the budget in problems 1-10.

The budget allocation task is illustrated in Figure 8. Subjects see one problem at a time, but they can scroll back and forth through the 10 problems until they are satisfied with their choices in all 10 problems. Subjects can also refer back to the instructions and tutorial until they are ready to submit their choices.

A.3.2 Budget allocation in problems 11-20

The method used for allocating the budget in these problems was case-by-case if the investment rule was used in problems 1-10, and investment rule if case-by-case was used in problems 1-10. The screenshots included here are for subjects selected to use the investment rule in problems 11-20.

As with case-by-case, the budget allocation task was preceded by instructions and a tutorial (Figures 9 and 10). The tutorial walks the subject through the use of the interface, using examples selected to illustrate its power. The investment rule tutorial includes nine tasks that have to be successfully completed before the next task is revealed. Figure 9 shows the tutorial before the first task is completed, and Figure 10 shows the tutorial after all nine tasks were completed. It is only after the successful completion of these nine tasks that the subject can proceed to allocate the budget in problems 11-20.

The budget allocation task is illustrated in Figures 11–13. The different options in the interface are explained in Section 2.2. Below the interface itself, subjects can see (i) a summary of the rule determined by the interface, and (ii) what the outcome of the rule would be on each of the 10 problems on which it would be implemented. Subjects can see one problem at a time, but they scroll back and forth through all 10 problems. A brief text below each

problem explains how the allocation in that problem follows from the rule that the subject has chosen. Figures 11–13 illustrate some of the investment rules that subjects could choose.

A.3.3 Choice of investment method for problems 21-30

In this part of the experiment, subjects had to decide which of the two investment methods to use for problems 21-30. As noted earlier, this part of the experiment is referred to as ‘Part 2’ in the paper, but in the instructions to subjects, it was referred to as ‘Part 3’. Figure 14 illustrates the choice screen for the $N = 180$ subjects whose rule option was to design a new rule for problems 21-30. The remaining $N = 144$ subjects were instead given the option to use the rule they designed in an earlier part of the experiment (either for problems 1-10 or for problems 11-20, as the case may be). Subjects in this second group were shown the outcome of this rule on problems 11-20 before making their decision (see also Section 2.5). The initial screen for this group is shown in Figure 15.

A.3.4 Budget allocations in problems 21-30

Subjects used the already familiar interface for their chosen investment method.

A.4 Optional survey

The survey (Figure 16) is optional. We ask subjects to complete it, but the wording (“We would be grateful if you answer this survey”) makes it clear that subjects can skip it with no financial penalty. The survey includes questions on both the case-by-case interface (Figure 17) and on the rule interface (Figure 18). Subjects are also asked about their decision in part 3. The questions differ depending on whether the subject chose the case-by-case (Figure 19) or the investment rule (Figure 20). Finally, subjects are asked demographic and other questions, some of which differ between Mturk and Lab subjects (Figure 21). Subjects are also asked about their decision in Part 2 (described as Part 3 in the experiment). The questions differ depending on whether the

subject chose the case-by-case interface (Figure 19) or the investment rule interface (Figure 20).

A.5 Results and payment

Finally, subjects were told what problem and company won, what allocation they made in that problem, and what their resulting payment would be. Subjects in mTurk conditions were told that they would be paid via the mTurk bonus mechanism, and lab subjects were asked to come to the lab administrator to collect their payment.¹⁹

¹⁹The screenshot is for an mTurk subject. Lab subjects had 2/3 the budget and were paid A\$1 for every 10 shares, rather than the US\$1 for every 100 shares of mTurk subjects.

General instructions

The experiment requires about 40 minutes of your time. Some parts can be demanding, and it is important that you are able to give it your full attention for the entire time. As in all experiments in economics, the procedures are described fully, and there is no deception. We appreciate your contribution to research!

The experiment includes 30 investment problems that share a common form. In each problem you have a budget of 600 Experimental Currency Units (ECUs) which you use to purchase shares in three companies that compete against each other.

After you complete the task, the computer will use a random number generator to select one of the 30 investment problems for payment, and to choose the winning company in that problem. Each of the 30 problems is equally likely to be selected for payment, and each of the three companies in that problem is equally likely to win.

Your investment in the winning company is worth real money to you. You will be paid one cent for every share you purchased in the winning company. This payment will be given to you as a bonus for this HIT. You will not be paid for any shares you purchased in the other companies.


The share prices of the three companies will be different in each investment problem you face. The price of shares does not affect your payment if the company wins nor its likelihood of winning, but it does affect the number of shares you can buy. Different people use different approaches in choosing their investment portfolio, and you should make the choices that suit you best.

Examples

The following two examples illustrate two different investment choices in the same problem. The price of shares in the three companies is, respectively, 1.00, 2.10, and 3.10 ECUs per share. In other problems the price of shares is different, but Company #1 is always the one with the lowest share price, and other companies are progressively more expensive. In all the problems the companies are listed in order, with Company #1 at the top, and Company #3 at the bottom.


Example 1

In this example, the entire budget has been invested in Company #1:



Example 2

In this example, the budget has been used to purchase the same number of shares in all companies:



Payment

Suppose the above decision problem has been selected for payment.

- If you were to make the investment choice in Example #1, you would receive a bonus of \$6.00 if Company #1 wins, and no bonus if some other company wins. Since there are three companies, your chance of receiving the \$6.00 bonus would be 1 out of 3, and your chance of receiving no bonus would be 2 out of 3.
- If you were to make the investment choice in Example #2, you would receive a bonus of \$0.97, whichever company wins. The reason is that you have 97 shares in **all** companies, so you are certain to have 97 shares in the winning company, no matter which one it's going to be.

Quiz

Before moving on, you are required to complete a quiz that tests your understanding of how your payment is determined. You can refer back to these instructions while completing the quiz.

Continue

Figure 4: General instructions (mTurk subjects in the 3 companies condition).

Quiz challenge #1

Suppose the problem below were selected for payment, and that you made the following investments:

price: 1.30	<input type="range" value="462"/>	462 shares
price: 1.60	<input type="range" value="0"/>	0 shares
price: 2.00	<input type="range" value="0"/>	0 shares
Budget:	<input type="range" value="0"/>	0 ECUs

What bonus would you receive if Company #1 is selected for payment?

\$ 4.62

What is the chance of this happening?

0%

1 in 3

2 in 3

100%

What bonus would you receive if some other company is selected for payment?

\$ 0

What is the chance of this happening?

0%

1 in 3

2 in 3

100%

Review instructions

Submit answer

Figure 5: The first quiz challenge (mTurk subjects in the 3 companies condition).

Quiz challenge #2

Suppose the problem below were selected for payment, and that you made the following investments:

price: 1.30	<input type="range" value="154"/>	154 shares
price: 1.60	<input type="range" value="125"/>	125 shares
price: 2.00	<input type="range" value="100"/>	100 shares
Budget:	<input type="range" value="0"/>	0 ECUs

What bonus would you receive if Company #1 is selected for payment ?

\$

What bonus would you receive if Company #2 is selected for payment ?

\$

What bonus would you receive if Company #3 is selected for payment ?

\$

Figure 6: The second quiz challenge (mTurk subjects in the 3 companies condition).

Part 1: Case-by-case investment decisions

In this part of the experiment, you will be allocating your budget in investment problems 1-10 by making **case-by-case investment decisions**. In each of these 10 investment problems, you will be making an independent decision how you want to allocate your budget in that particular problem.

Tutorial

This tutorial uses a number of simple tasks to explain the case-by-case investment decisions interface. You are required to complete the tutorial before you can start making your investment decisions in problems 1-10. At each point in the tutorial, you should complete the task that is highlighted in green to reveal the next task. All the tasks refer to the sample problem at the bottom of the page.

1. You allocate your budget by dragging the sliders of the different companies. Start by pushing the slider of Company #3 to the right as far as it goes.
2. Notice that the sliders of the other companies move together with the slider of Company #3, so that you end up with the same number of shares in all three companies. This is convenient, and keeps you from accidentally buying more shares in Company #3 than in any of the other companies. Since all companies have the same chance of winning, you should always buy at least as many shares in Company #1 as in Company #2, and at least as many shares in Company #2 as in Company #3. Now drag the slider of Company #1 back to zero.
3. Dragging the slider of Company #1 back to zero resets the sliders of all companies, and is a quick way of starting afresh. Now push the slider of Company #1 to the right as far as it goes.
4. You have now allocated your entire budget to Company #1. In this example, the price of shares in Company #1 is 1.00, so you were able to drag the slider all the way to right, and purchase a total of 600 shares. This is not always the case. In problems in which the price of shares in Company #1 is more than 1.00, you can only drag the slider part of the way before your budget is exhausted, and the number of shares you end up with is less than 600 shares.

There are many other ways for allocating your budget. Your final task in this tutorial is as follows. You should allocate all your budget, purchasing **some** shares in all companies, and overall purchasing **more** shares in Company #1 than in Company #2, and **more** shares in Company #2 than in Company #3.

Congratulations! You have now finished the tutorial, and can proceed to allocate your budget in investment problems 1-10. In order to do so, click the 'Continue' button at the bottom of the page.

Tutorial sample problem

price: 1.00	<input type="range" value="425"/>	425 shares
price: 2.10	<input type="range" value="51"/>	51 shares
price: 3.10	<input type="range" value="22"/>	22 shares
Budget:	<input type="range" value="0"/>	0 ECUs

Figure 7: The case-by-case instructions and tutorial, shown here for a subject who has already successfully completed the four tutorial tasks.

Investment problems 1-10

Your task is to make case-by-case investment decisions in problems 1-10. Use the **◀** and **▶** arrows to move between problems. Submit your choices once you are happy with your budget allocations in all the problems.

◀ Investment problem #1 ▶

price: 1.10	<input type="range" value="155"/>	155 shares
price: 1.30	<input type="range" value="155"/>	155 shares
price: 3.00	<input type="range" value="0"/>	0 shares
Budget:	<input type="range" value="228"/>	228 ECUs

[Review instructions ▲](#) [Submit your choices](#)

Figure 8: The case-by-case allocation screen.

Part 2: Investment rule

In this part of the experiment, you will be allocating your budget in investment problems 11-20 with the help of an **investment rule**. The investment rule is a systematic approach for allocating the budget in different problems. The investment rule you choose will be used to allocate your budget in all 10 problems.

Tutorial

The tutorial uses a number of simple tasks to explain the different options you can choose from. The tasks refer to the investment rule interface below the tutorial. At each point in the tutorial, you should complete the task that is highlighted in green to reveal the next task. When you complete the final task you will be able to choose your investment rule for problems 11-20.

1. **The simplest investment rule is to invest your entire budget in Company #1.** *Choose this rule now by selecting "Company #1 (only)" in the first drop-down list under the heading "Which companies do you want to invest in?"*

Note: the "baseline number of shares" checkbox should remain clear at this point

The investment rule interface

Which companies do you want to invest in?

(Select) ▾

Depending on the difference in share price: invest in all the companies that are at most (Select) ▾ more expensive than Company #1.

How do you want to split your budget between these companies?

Purchase the same number of shares in each of these companies

Invest the same number of ECUs in each of these companies

Do you want to purchase a baseline number of shares in all companies?

Yes: purchase a baseline of (Select) ▾ shares in all companies

Investment rule summary: Your investment rule is incomplete. You must choose which companies you want to invest in.

Continue

Figure 9: The investment rule tutorial. This screen shows the initial screen before the first tutorial task was completed.

1. The simplest investment rule is to invest your entire budget in Company #1. Choose this rule now by selecting "Company #1 (only)" in the first drop-down list under the heading "Which companies do you want to invest in?"
2. Below the rule interface itself you can now see the investment rule summary, the outcome of this investment rule in three sample investment problems, and an explanation for how the rule works in each of these problems. Visit the other two sample problems now to see the outcome of the rule in these problems.
3. Allocating the entire budget to Company #1 buys the greatest possible number of shares, but you only get a bonus if Company #1 is the winning company. The opposite approach is to invest in all companies. Choose this rule now, by changing your selection in the first drop-down list to the final option: "All three companies"
4. If you choose to invest in more than one company, you have two options for splitting your budget between these companies. You can either purchase the same number of **shares** in all companies, or you invest the same number of **ECUs** in all companies. Try the first option now by selecting "purchase the same number of shares in each of these companies" in the second part of the rule.
5. If you look at the sample problems, you can see that with this rule you always end up with the same number of shares in all the companies. Now try the other option by selecting "Invest the same number of ECUs in each of these companies" in the second part of the rule.
6. Because the price of shares differs between companies, this rule purchases more shares in Company #1 than in Company #2, and more shares in Company #2 than in Company #3. Visit the other two sample problems now to see the outcome of the rule in these problems.
7. If you prefer, you can let the number of companies you invest in differ between problems, depending on the price of shares in these problems. For example, you can invest in all the companies that are at most twice as expensive as Company #1. Try this now by choosing a threshold of 100% in the option that says 'Depending on the difference in share price'.
8. In order to appreciate how this option works, you should view the outcome in the three sample problems. Visit the other two sample problems now to see the outcome of the rule in these problems.
9. The third and final part of the investment rule enables you to purchase a baseline number of shares in all companies. For example, you can purchase 25 shares in all companies, and invest the rest of your budget in Company #1. Try this now, by selecting "Company #1 (only)" in the first part of the rule, and 25 in the drop-down list for "baseline number of shares".

Figure 10: The investment rule tutorial. This screen shows the screen after all the tutorial tasks were successfully completed.

Which companies do you want to invest in?

Company #1 (only) ▾

Depending on the difference in share price: invest in all the companies that are at most (Select) ▾ more expensive than Company #1.

How do you want to split your budget between these companies?

Purchase the same number of shares in each of these companies


Invest the same number of ECUs in each of these companies

Do you want to purchase a baseline number of shares in all companies?

Yes: purchase a baseline of 50 ▾ shares in all companies

Investment rule summary: purchase a baseline of 50 shares in all companies, and invest the rest of your budget in Company #1.

The panel below shows the outcome of your investment rule in each of the 10 investment problems.



Explanation: purchasing the baseline of 50 shares in all companies cost 53% of your budget. The rest was invested entirely in Company #1, purchasing an additional 259 shares in Company #1.

Figure 11: The investment rule interface: baseline option example. This screenshot shows the impact of the rule on problem #13 in a three company condition. Subjects can view the impact of the rule on all the 10 problems on which it is implemented.

Which companies do you want to invest in?

(Select) ▾

Depending on the difference in share price: invest in all the companies that are at most more expensive than Company #1.

How do you want to split your budget between these companies?

Purchase the same number of shares in each of these companies


Invest the same number of ECUs in each of these companies

Do you want to purchase a baseline number of shares in all companies?

Yes: purchase a baseline of shares in all companies

Investment rule summary: use your budget to purchase the same number of shares in Company #1 and in any other companies that are at most 100% more expensive than Company #1.

The panel below shows the outcome of your investment rule in each of the 10 investment problems.



Explanation: since the price of shares in Company #2 is within 100% of the price of shares in Company #1, and the price of shares in Company #3 is above this threshold, your budget was used to purchase 207 shares in companies #1 and #2.

Figure 12: The investment rule interface: threshold rule example. This screenshot shows the impact of the rule on problem #13 in a 3 company condition. Subjects can view the impact of the rule on all the 10 problems on which it is implemented.

Which companies do you want to invest in?

All three companies ▾

Depending on the difference in share price: invest in all the companies that are at most (Select) ▾ more expensive than Company #1.

How do you want to split your budget between these companies?

Purchase the same number of shares in each of these companies


Invest the same number of ECUs in each of these companies

Do you want to purchase a baseline number of shares in all companies?

Yes: purchase a baseline of (Select) ▾ shares in all companies

Investment rule summary: split your budget equally between all three companies.

The panel below shows the outcome of your investment rule in each of the 10 investment problems.



Explanation: your budget was split equally between the three companies, purchasing 182 shares in Company #1, 111 shares in Company #2, and 59 shares in Company #3.

Figure 13: The investment rule interface: equal number of ECUs in each company example. This screenshot shows the impact of the rule on problem #13 in a 3 company condition. Subjects can view the impact of the rule on all the 10 problems on which it is implemented.

Part 3

In this third and final part of the experiment, you will be allocating your budget in investment problems 21-30. You can choose the method you prefer.

- If you choose **case-by-case investment decisions**, you would be using the interface from the first part of the experiment to make case-by-case decisions in these problems.
- If you choose **investment rule**, you would be using the investment rule interface from the second part of the experiment to choose an investment rule for these problems.

Submit your choice ▾

Case-by-case investment decisions

Investment rule

Figure 14: Decision screen in part 3.

Part 3

In this third and final part of the experiment, you will be allocating your budget in investment problems 21-30. You can choose the method you prefer.

- If you choose **investment rule**, the investment rule you chose in the first part of the experiment would be applied to these problems.
- If you choose **case-by-case investment decisions**, you would be using the interface from the second part of the experiment to make case-by-case decisions in these problems.

In order to assist you in making this decision, you can review your investment rule from the first part of the experiment, and see the resulting budget allocations in problems 21-30.

Continue

Figure 15: Part 3 decision in the older design

Survey

You have now completed all the tasks in experiment. We would be grateful if you answer this survey before proceeding to the results. Doing so would help us analyze the results of the experiment, and improve the experiment in the future.

Figure 16: The start of the survey.

Case-by-case investment decisions

How did you decide on your investment decisions in the different problems?

Did you have a good understanding of the case-by-case investment decisions interface?
No Yes
Were you able to use the interface to allocate your budget the way you wanted to?
No Yes
What, if anything, did you want to do, but weren't able to?

Figure 17: Survey questions on the case-by-case interface.

Investment rule

How did you decide on your investment rule?

Did you have a good understanding of the investment rule interface?
No Yes

Were you able to use the interface to allocate your budget the way you wanted to?
No Yes

What, if anything, did you want to do, but weren't able to?

Figure 18: Survey questions on the investment rule interface.

Your preference between the two. In part 3 of the experiment, you chose case-by-case investment decisions over the investment rule

How would you characterize your certainty in this choice?
weak preference strong preference

Which of the following best characterizes your reason for preferring the case-by-case interface?

I couldn't choose the investments I wanted with the investment rule

I just like having direct control

Some other reason

Anything else?

Which way of allocating your budget fits better the way you think?
case-by-case investment decisions investment rule

Figure 19: Survey questions on part 3 choice for subjects who chose to make case-by-case choices in problems 21-30. Compare with Figure 20.

Your preference between the two. In part 3 of the experiment, you chose the investment rule over case-by-case investment decisions

How would you characterize your certainty in this choice?
weak preference strong preference
_____○_____

Which of the following best characterizes your reason for preferring the investment rule?

- The investment rule allowed me to invest the way I wanted to
- Using the case-by-case interface would have taken longer
- Some other reason

Anything else?

Which way of allocating your budget fits better the way you think?
case-by-case investment decisions investment rule
_____○_____

Figure 20: Survey questions on part 3 choice for subjects who chose to use to design a new rule for problems 21-30. Compare with Figure 19.

Other questions

Would you describe yourself as someone who tries to avoid risks whenever possible or as someone who is comfortable taking risks?
avoids risks whenever possible comfortable taking risks
_____○_____

Gender
 Female Male

Are you a native speaker of English?
 No Yes

Year of birth

What course do you study?

Figure 21: Other survey questions (lab version).