

NBER WORKING PAPER SERIES

COMPLEXITY AND SOPHISTICATION

Leandro Carvalho
Dan Silverman

Working Paper 26036
<http://www.nber.org/papers/w26036>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2019, Revised February 2022

This paper benefited from many discussions with Shachar Kariv, and participants in several seminars and conferences. Special thanks to Alisdair McKay, Andrew Caplin, Juan Carillo, Mark Dean, Alex Rees-Jones, John Leahy, Dmitry Taubinsky, and Jeremy Tobacman for their thoughtful comments, and to Carla Blair, Tania Gutsche, Adrian Montero, Bart Oriens, and Bas Weerman for their integral work on implementing the experiment. Francesco Agostinelli provided exemplary research assistance. This work was funded by the National Institute on Aging (NIA P30AG024962) through USC's Roybal Center for Financial Decision Making. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Leandro Carvalho and Dan Silverman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Complexity and Sophistication
Leandro Carvalho and Dan Silverman
NBER Working Paper No. 26036
July 2019, Revised February 2022
JEL No. D14,D81,G02,G11

ABSTRACT

Many financial situations present individuals with simple alternatives to solving complex problems. The value of these alternatives depends on whether individuals are sophisticated and know when they are better off opting out of complexity. We tested complexity's effects and evaluated sophistication in a large and diverse sample. We randomly assigned both complexity to portfolio problems and the offer of a simple alternative to portfolio choice. The less-skilled earn lower returns under complexity because they often opt out when the simple alternative is dominated. To interpret the findings, we develop a novel method for estimating the structural parameters of a rational inattention model. The structural estimates are consistent with substantial sophistication. Large fractions of the money that the low-skilled lose by opting out can be justified by their higher costs of attention.

Leandro Carvalho
Center for Economic and Social Research
University of Southern California
635 Downey Way
Los Angeles, CA 90089
leandro.carvalho@usc.edu

Dan Silverman
Department of Economics
W.P. Carey School of Business
Arizona State University
P.O. Box 879801
Tempe, AZ 85287-9801
and NBER
Daniel.Silverman.1@asu.edu

1 Introduction

As financial markets evolve and the set of financial instruments grows larger, individuals are asked to make saving, credit, and insurance choices in an increasingly complex environment. New alternatives can improve individual welfare, but the additional complexity likely makes optimization more difficult, and may thus reduce the quality of financial decisions. These pitfalls of complexity might be avoided at low cost, however, if individuals are sophisticated and know when they should choose simple alternatives rather than solve complex problems. If, for example, a worker knows he will struggle to choose and then properly adjust his retirement savings portfolio, and if he feels confident in his firm's default portfolio allocation, he can accept the default and avoid both the costs of considering and managing all his options, and the risk of making a badly suboptimal choice.

The importance of simple alternatives to solving complex financial problems, and thus sophistication, is rising. In 2018, for example, 52% of all participants in Vanguard employer-sponsored savings plans were investing solely in a single target-date retirement fund; in the Australian pension system, 64% of total contributions to superannuation funds in 2015 were invested in default portfolios chosen by the employer.¹

This paper presents the results of an experiment to study the effects of complexity on financial choices and to evaluate the sophistication of individuals to know when they are better off taking a simple option instead of solving a complex problem. The experiment involved a diverse sample of 700 U.S. participants who each made many incentivized investment portfolio choices. The complexity of the investment problems was randomly assigned, and determined by the number of assets in which the participant could invest. Importantly, as the number of assets changed the real investment opportunities did not. The additional assets did not replicate those in the simple problem, but they were redundant; any distribution of payoffs that was feasible in a simple problem was also feasible in a complex problem, and vice versa. We therefore interpret the treatment as isolating the influence of complexity separate from other effects of adding options to an opportunity set.

Participants were also randomly assigned the opportunity to take a deterministic outside option rather than make an active portfolio choice. The payoff from the outside option varied randomly and was sometimes dominated by the payoff from the riskless portfolio in the investment problem, i.e. the asset allocation with a deterministic return. These outside options were meant to capture investment opportunities, such as default portfolios, target-date retirement savings plans, or age-based college saving plans, that require less consideration or management on the part of the individual, but may not be well-tailored to her objectives.

The results show that, when they are required to make an active portfolio decision, participants facing complex problems choose allocations with lower expected returns and lower risk. These

¹“How America Saves 2019,” Vanguard Research Group. Australian Prudential Regulation Authority “Annual Superannuation Bulletin,” August 2016.

effects of complexity on choices are not attributable just to changes in well-behaved risk preferences. Instead we find evidence of a decline in decision-making quality as measured by violations of normative choice axioms (cf. Choi et al. 2014). In particular, complexity leads to statistically significant increases in violations of dominance principles including monotonicity with respect to first-order stochastic dominance.

The ability to opt out does not protect participants against the negative effects of complexity. Instead, the availability of the outside option has a substantial negative effect on expected payoffs. This effect is larger for those with the lowest levels of financial decision-making skills, defined in terms of numeracy, financial literacy, and consistency with utility maximization in another experiment. The option to avoid complexity reduces their expected payoff by 15 percentage points. These declines in expected payoffs are almost exclusively due to the choice by the least-skilled to opt out when the outside option is dominated, and not to the effects of having the option (and not taking it) on their portfolio choices.

While low-skilled participants often earn sharply lower returns when they opt out, their decision to avoid complex portfolio problems may nevertheless be sophisticated. They may be better off avoiding the costs of attending to a portfolio problem even if that implies foregoing high-return investment opportunities. Indeed, metadata on participants' choices are consistent with the idea that the low-skilled face higher costs of attending to complex problems. We find that, even when they have no opportunity to opt out of complexity, low-skilled participants spend substantially less time making their choices and make fewer active adjustments to their portfolios. This suggests they are paying less attention to the problems, perhaps because they face higher costs of doing so.

It is challenging to evaluate the sophistication of decisions to opt out because a central determinant of that evaluation, attention costs, are unobserved. We turn to a rational inattention model to aid interpretation and draw inference on these unobserved costs. That model interprets differences in behavior as resulting either from differences in the payoffs associated with each available option, or from differences in beliefs about the payoff from each option, or from differences in the cost of attending to the payoff-relevant features of the problem. In this way, the model treats all behavior as rational. We interpret higher rates of dominated choices as “sophisticated” to the extent they can be explained by higher costs of attending to the portfolio problem.

To obtain structural estimates of heterogeneous attention costs without making strong assumptions about participants' utility functions or prior beliefs, we develop a novel identification strategy that may be applied in other settings. That strategy exploits the random variation in experimental endowments and outside options to difference out heterogeneity in preferences and beliefs. This strategy also avoids making strong assumptions on how the discrete space of actions is defined. Intuitively, the identification strategy approximates an experiment where treatment and control, who face the same asset prices and the same returns, are assigned different endowments or outside options. Effectively holding individuals' preferences and beliefs fixed, attention costs are identified

by the sensitivity of the conditional probability of choosing the outside option to changes in the payoff from that option and to changes in the experimental endowment. We show that utility need not be linear to draw this inference; it is sufficient that it be homothetic.

The estimates of attention costs derived via this identification strategy fail to reject a null hypothesis of sophistication among the low-skilled. Point estimates indicate that, relative to the higher-skilled, low-skilled participants have a significantly higher cost of acquiring information about payoffs under complexity. The low-skilled are, that is, less responsive to the relative return of working through complex decision problems. Simulation of the estimated model shows, moreover, that the implied costs of attention are sufficient to explain a large fraction of the money that the low skilled leave on the table by making dominated choices.

1.1 Related Literature

This paper joins an economics literature on the influence of complexity and the problem of evaluating large menus of choices. That literature includes several theories of complexity and models of choice from large sets. See, for example, Wilcox (1993); Al-Najjar et al. (2003); Gale and Sabourian (2005); Masatlioglu et al. (2012); Ortoleva (2013); and Caplin and Dean (2015). These theories are motivated by common sense and by a long tradition (cf. Simon 1957) of accounting for decision-makers' costs of attending to all feasible options.

Interest in complexity and the problems of large choice sets is also motivated by a substantial experimental literature focused on the effects of increasing the number of alternatives from which a decision-maker may choose.² Iyengar and Lepper's (2000) influential field experiment in a grocery store provided evidence of a "paradox of choice," where having too many options (of jam) may demotivate buying.³ Related studies have examined the effects of a larger number of options on portfolio choices (Agnew and Szykman, 2005; Iyengar and Kamenica, 2010; Beshears et al., 2013), procrastination (Tversky and Shafir, 1992; Iyengar, Huberman and Jiang, 2004), and status quo bias (Samuelson and Zeckhauser, 1988; Kempf and Ruenzi, 2006; Dean, 2008; Ren, 2014). A common feature of these studies is that the opportunity set changes across the simple and complex conditions. This feature captures an important aspect of complexity in reality, but it may confound the influence of complexity with more or less standard effects of a larger choice set.

Motivation for studying complexity also comes from influential field evidence of the effects of retirement saving plan defaults in a literature sparked by Madrian and Shea (2001). One explanation for the important effects of defaults is that employees interpret default status as an endorsement of a plan by their employer. Accepting the default can thus reflect a sophisticated response to the

²See Tse et al. (2014); Friesen and Earl (2015); Abeler and Jager (2015) for examples of other dimensions of complexity that have been studied.

³There are, however, many studies that find no such effect of increasing the number of choices on a menu (Scheibehenne, Greifeneder and Todd, 2010).

costs of attending to complex financial decisions (Beshears et al., 2009). While we study active choices rather than defaults,⁴ the experiment provides systematic evidence on the plausibility of sophisticated decisions to opt out of complex financial decisions.

The present paper also contributes to a literature on the effects of more options on the quality of decision-making.⁵ Using designs where some (sets of) choices may violate normative axioms, a few studies find that complexity reduces the likelihood of making good choices (Caplin, Dean, and Martin, 2011; Schram and Sonnemans, 2011; Besedes et al., 2012a; Brocas et al., 2014; Kalayci and Serra-Garcia, 2015).⁶ Similarly, Carlin, Kogan, and Lowery (2013) find that complexity in asset trading leads to increased price volatility, lower liquidity, and decreased trade efficiency.

This paper advances the frontier of research in these areas in three main ways. First, by keeping real opportunity sets constant across treatments, the experiment separates the influence of complexity on financial choices from other effects of increasing the number of options in a menu. Second, by implementing the experiment with a web-based panel, the experiment studies these effects of complexity on financial choices in a large and diverse sample about which much is already known. The size, heterogeneity, and existing measures of the sample allow disaggregated study and evaluation of external relevance.

Last, this paper advances a novel interpretation of, and methods for studying, complexity's effects on choice, including the choice to opt out. The interpretation is centered on rational inattention and sophistication, and the methods involve new and robust techniques for estimating attention costs. Prior empirical work on complexity has demonstrated a variety of effects but has not provided quantitative interpretations with a structural model of rational inattention, or studied the sophistication of decisions to opt out. The closest economics research on this form of sophistication appears in studies of the demand for and consequences of professional advice, but we are not aware of prior research that uses the structural estimates of a rational inattention model for this purpose.⁷ The method is distinctive as it allows point estimates of heterogeneous costs of attention without having to simultaneously estimate the prior beliefs about the optimality of each feasible choice or assume these prior beliefs are uniform. The method also requires only limited assumptions on how the discrete action space is specified. In this way, the paper offers a new, portable, and robust

⁴For each participant, and for each problem in the experiment, the initial asset allocation was chosen at random. If defaulting is interpreted as actively choosing this initial allocation, then participants default on 9.5% of all choices. The default rate under complexity is lower, 7.7%.

⁵Huck and Weizsacker (1999) show complexity reduces the likelihood that participants maximize expected value.

⁶One exception is Besedes et al. (2012b).

⁷See, e.g., Chalmers and Reuter (2012), Mullainathan, et al. (2012), and Egan et al. (2016). Chalmers and Reuter (2012), for example, find evidence that when the Oregon University System stopped offering retirement savings plan participants free access to advice from brokers, many new participants were more likely to take up target date retirement funds and arrive at portfolios that were superior in terms of risk and return. The closest use of structural models to draw inference about choice and belief imperfections may be those concerned with tax salience and misperceptions. Examples of these studies include Chetty, Looney, and Kroft (2007), Taubinsky and Rees-Jones (2016), and Rees-Jones and Taubinsky (2016).

strategy for estimating key parameters of a rational inattention model.

2 Study Design

In this section, we present a conceptual framework for the study and then describe the experimental procedures.

2.1 Conceptual Framework for the Experiment

The literature offers several definitions of complexity. Here we adopt a definition based on Al-Najjar et al. (2003) that ranks the complexity of any two choice problems X and Y that are sufficiently familiar to the decision-maker. For our purposes, sufficient familiarity means that the decision maker understands that optimal choice in these problems depends on a finite set of contingencies. The decision-maker’s preferred option is contingent, that is, on the realization of a discrete random variable that determines the underlying value of each available option.

Definition 1 *A decision maker perceives choice problem X to be **more complex** than Y if she thinks optimization over the choices in X requires consideration of more contingencies than optimization over the choices in Y .*

To isolate the effects of this notion of complexity on decision-making, we designed two problems – one simple, one complex – that share the same opportunity set. In the two problems, participants are given an endowment to invest in risky assets. The assets have different prices, and different payouts that depend on whether a coin comes up heads or tails. The only distinction between the simple and the complex problems is that in the simple problem there are two assets while in the complex problem there are five.

Figure 1: Illustration of Simple and Complex Problems

Simple Problem			Complex Problem				
	A	B	A	B	C	D	E
Prices	\$0.90	\$1.00	\$0.90	\$1.00	\$0.93	\$0.96	\$0.99
Payouts							
Heads	\$0.00	\$2.00	\$0.00	\$2.00	\$0.60	\$1.20	\$1.80
Tails	\$2.00	\$0.00	\$2.00	\$0.00	\$1.40	\$0.80	\$0.20

Figure 1 illustrates with an example. In the simple problem there are two investment options: assets A and B. Each share of asset A has a price of \$0.90 and each share of asset B has a price of \$1. Each share of asset A pays \$0 in the case of heads and \$2 if tails. Asset B shares pay \$2 if heads and \$0 if tails. The options in the complex problem include the two assets available in the

simple problem plus three additional assets – C, D, and E – each of which is a convex combination of assets A and B. Asset C is composed of 70% asset A and 30% asset B; Asset D is 40% of asset A and 60% asset B; and asset E is 10% asset A and 90% asset B. Because assets C, D, and E are convex combinations of assets A and B, any portfolio in the complex problem can be re-created in the simple problem, and vice versa (see Online Appendix for a proof).⁸

According to Definition (1), the five-asset problem is more complex than the two-asset problem if the participant thinks optimal choice requires consideration of more contingencies in the problem with more assets. This is natural if the interpretation of contingencies includes the relative prices of the assets or their returns.

2.2 Sample

The study was conducted with 700 members of the University of Southern California’s Understanding America Study (UAS), an Internet panel with respondents ages 18 and older living in the U.S. Respondents are recruited by address-based sampling. Those without Internet access at the time of recruitment are provided tablets and Internet access. About twice a month, respondents receive an email with a request to visit the UAS site and complete questionnaires.

The study consisted of one baseline and one follow-up survey. In the baseline survey participants were administered Choi et al.’s (2014) choice under risk experiment. As explained below, these choices can be used to construct baseline measures of decision-making skills. In the follow-up survey we administered a collection of the simple and complex problems described above. In addition, panel members provided a variety of information collected in previous UAS modules. This information includes demographics and socioeconomic data, and results of numeracy and financial literacy tests.

2.3 Experimental Design

The experiment had a 2 x 2 between-subjects design, where participants were randomly assigned to one of four treatment arms as shown in the table below. One manipulation involved varying the number of investment options: Arms I and II were assigned to the simple problem with two assets while arms III and IV were assigned to the complex problem with five assets.⁹

⁸By using convex combinations of assets to span the same opportunity set, our design resembles that in Eyster and Weizsäcker (2010). In that paper, participants faced a sequence of portfolio problems, each with two assets. Each problem had a twin, with the same opportunity set but formed by assets with differently correlated returns. The paper’s basic question was whether participants neglected that difference in correlation.

⁹A participant saw either the simple or the complex framing of problems, but not both. This would seem to make it less likely that a participant assigned to the complex frame could see the redundancy of the assets. Indeed, participants assigned to treatment arm III invested exclusively in 2 Arrow securities only in 14.1% of their decisions (in 7.0% of their decisions participants invested the entire endowment in just one of the Arrow securities).

	Simple Problem	Complex Problem
Forced to Invest	I	III
Option to Avoid Investment	II	IV

The other manipulation involved offering participants the option of avoiding the investment problem. Participants assigned to arms II and IV were offered the choice between making the investment decision or taking an “outside option” of \$2, \$5, \$10, \$15, or \$20. The amount of the outside option was randomly varied across participants.

2.4 Experimental Task

The experimental task involved participants allocating their endowment across two (treatment arms I and II) or five (treatment arms III and IV) assets. They were given information about the price per share of the assets and how much the assets paid depending on the coin toss. Participants made their investment choices by choosing the number of shares of each asset they wanted to buy.

Online Appendix Figure 1 shows a screenshot of the interface used in treatment arms I and II. The table at the top of the screen shows the prices of assets A and B and their payouts. The participant was then shown the amount available for investing and prompted to make her investment choices. The graph below the table displays the number of shares owned of each asset. Participants made their allocations by either dragging the bars or by clicking the + and – buttons.¹⁰

Treatment arms III and IV used a similar interface (see Online Appendix Figure 2). The distinction is that they were shown information about 5 assets – A, B, C, D, and E – and the graph displayed 5 bars. Participants were shown a tutorial video to learn how to use the interface and had two rounds to practice – participants assigned to the simple and complex conditions were shown the same tutorial video and were administered the same practice trials; in both the tutorial video and in the practice trials the endowment could be invested in 3 assets.¹¹ We randomized the initial levels of the bars.

The interfaces for treatment arms II and IV were different because these groups were offered the option to avoid the investment decision. Online Appendix Figure 3 shows a screenshot of the interface for treatment arm II. It differs from the interface for treatment arm I in two ways. First, the graph with the bars is not shown. Second, the sentence “How many shares of each asset do you want to buy?” is replaced by a prompt for the participant to choose between investing the experimental endowment (button “Invest \$26”) and taking the outside option (button “Receive

¹⁰The interface was such that participants always invested 100% of their experimental endowment.

¹¹<https://www.youtube.com/watch?v=TNr3Wgakczk&feature=youtu.be> We conducted cognitive interviews to ensure participants understood the tutorial video and what they were supposed to do in the experiment.

\$5”). If she clicked on the first button, the bars were unveiled and she could make her investment choices using the same interface used by treatment arm I. If she clicked on the second button, she was presented with the next problem.

Participants were presented with 25 investment problems (one was randomly selected for payment; the participant was paid the outside option if in the problem selected for payment she chose to avoid). They were not given feedback during the experiment.

It is useful to think of each problem as a two-dimensional budget line, where the axes correspond to payoffs in two states of the world: heads (y-axis) and tails (x-axis). The y-intercept is the payoff if the endowment is invested all in heads and the coin comes up heads; the x-intercept is the payoff if the endowment is invested all in tails and the coin comes up tails. The investment problems were drawn at random from 10 sets, each of which consisted of 25 budget lines. In each set, lines were chosen to generate substantial variation in the relative prices of the assets and in the endowment available for investment.¹² The order in which the budget lines were presented to each participant was also randomized.

Each budget line was then converted into a simple problem using the following procedure. Let asset A be the asset that pays \$2 if the coin comes up tails and \$0 otherwise and let asset B be the asset that pays \$2 if the coin comes up heads and \$0 otherwise. We normalized the price of asset B to \$1 such that the endowment was equal to the y-axis intercept divided by 2 (rounded to closest integer for convenience). The price of asset A was equal to the y-axis intercept divided by the x-axis intercept (rounded to closest multiple of 0.1). For example, a budget line with a x-intercept of \$40 and a y-intercept of \$80 would be converted into a simple problem with an endowment of \$40 and where each share of asset A would cost \$2. We randomized the order in which assets A and B were shown on the screen. As described above, to construct a complex analogue of a simple problem we created assets C, D, and E by taking convex combinations of the prices and payouts of assets A and B. We randomized the order in which assets A, B, C, D, and E were shown, from left to right, on the screen.

2.5 Measuring the Quality of Decision-making

We exploit the within-subject variation in the endowment and in asset prices to construct individual-specific measures of decision-making quality. We examine four measures of the quality of decision-making. First, we study whether choices violate the General Axiom of Revealed Preference (GARP). Choi et al. (2014) and Kariv and Silverman (2013) argue that consistency with GARP is a necessary but not sufficient condition for high-quality decision-making. This view draws on

¹²We used a procedure similar to the one used by Choi et al. (2014) to draw budget lines. First we randomly selected between the x-axis and y-axis. If the y-axis was selected, we chose the y-intercept by drawing uniformly between \$10 and \$100. If the chosen y-intercept was greater than \$50, we would draw the x-intercept uniformly between \$10 and \$100. If the selected y-intercept was smaller than \$50, we would draw the x-intercept uniformly between \$50 and \$100. We dropped budget lines where $y\text{-intercept} < 0.05 * x\text{-axis intercept}$.

Afriat (1967), which shows that if an individual’s choices satisfy GARP in a setting like the one we study, then those choices can be rationalized by a well-behaved utility function. Consistency with GARP thus implies that the choices can be reconciled with a single, stable objective. Here we will assess how nearly individual choice behavior complies with GARP using Afriat’s (1972) Critical Cost Efficiency Index (CCEI). The CCEI is a number between zero and one, where one indicates perfect consistency with GARP. The degree to which the index falls below one may be viewed as a measure of the severity of the GARP violations.¹³

Consistency with GARP may be viewed as too low a standard of decision-making quality because it treats all stable objectives of choice as equally high quality.¹⁴ A more stringent requirement would also require monotonic preferences and, because the realization of the state (heads or tails) should not influence the utility from money, symmetry of demand for these assets. In particular, violations of monotonicity with respect to first-order stochastic dominance (FOSD) – that is, the failure to recognize that some allocations yield payoff distributions with unambiguously lower returns – may be regarded as imperfect choices and provide a more compelling criterion for high-quality decision-making. Similarly, asymmetries of demand with respect to the state of the world might also be regarded as evidence of lower decision-making quality.

To illustrate a violation of first-order stochastic dominance, suppose that the y-axis intercept is larger than the x-axis intercept (such that the price of tails is higher than the price of heads) and that a participant chooses an allocation (x,y) that is on the “shorter side” of the 45 degree line. It is possible to show that there is an allocation (y,z) on the “longer side” of the 45 degree line that yields an unambiguously higher payoff distribution than (x,y) – i.e., $z > x$. The third measure of decision-making quality is the fraction of times in which participants selected a dominated portfolio.¹⁵

Following Choi et al. (2014), we calculated a FOSD score as follows. If there was no feasible allocation that dominated the selected allocation, then the FOSD score was assigned a value of 1. If the selected allocation was dominated, we calculated the FOSD score as $\frac{x+y}{z+y}$, which equals the expected return of the selected allocation as a fraction of the maximal expected return. For participants assigned to treatment arms II and IV, we used the same procedure to calculate the

¹³Formally, the CCEI measures the fraction by which all budget lines described above must be shifted in order to remove *all* violations of GARP. Put precisely, suppose the choice data for individual i are given by $\mathbf{p}^i, \mathbf{x}^i$ where the vector \mathbf{p}^i describes the relative prices (budget sets) i faced, and \mathbf{x}^i describes the choices made from those budget sets. Then for any number $0 \leq e \leq 1$, define the direct revealed preference relation

$$\mathbf{x}^i R^D(e) \mathbf{x}^j \Leftrightarrow e \mathbf{p}^i \cdot \mathbf{x}^i \geq \mathbf{p}^i \cdot \mathbf{x}^j,$$

and define $R(e)$ to be the transitive closure of $R^D(e)$. Let e^* be the largest value of e such that the relation $R(e)$ satisfies GARP. The CCEI is the e^* associated with the data set.

¹⁴For example, consider a participant that always allocates her endowment to heads. This behavior is consistent with maximizing the utility function $U(x_{heads}, x_{tails}) = x_{heads}$ and would generate a CCEI score of one. However, these choices are hard to justify because for some of the budget lines, allocating the endowment to heads means allocating it to the more expensive asset, a violation of monotonicity with respect to first-order stochastic dominance.

¹⁵We drop choice sets where the price of asset 1 equals \$1 and thus all portfolios have the same expected return.

FOSD when participants chose to make investment decisions. However, when they chose to avoid decision-making, we calculated the FOSD score as $\min \left\{ 1, \frac{\text{outside option}}{\text{risk free return}} \right\}$.

To provide a unified measure of violations of GARP, of monotonicity with respect to FOSD, and symmetry of demand, we combine the 25 choices for a given participant with the mirror image of these data obtained by reversing the prices for heads and tails and the actual choices. More specifically, if (x_1, x_2) were chosen from the budget defined by $(p_1, p_2; m)$ where $p_1x_1 + p_2x_2 = m$, then we assume (x_2, x_1) would have been chosen subject to the mirror-image budget $(p_2, p_1; m)$. We then compute the CCEI for the data set consisting of 50 observations that combines the 25 actual choices with their 25 mirror images. Violations of GARP in this combined dataset result from violations in the actual choice data and from violations of monotonicity with respect to FOSD. This measure can also detect violations of symmetry, as the combined data would induce no more violations if the actual demand were consistent with GARP and symmetric. Cf. Choi et al. (2014).

3 Descriptive Results

3.1 Summary Statistics

Summary statistics of the sample show that the controls are balanced across treatment arms. The first four columns of Table 1 show means, separately by treatment arm. Participants ranged in age from 18 to 90 with an average and median age of 48. There is also substantial variation in schooling (21% had a high school diploma or less while 57% graduated from college) and in annual household income (with 25% making \$30,000 or less and 20% making \$100,000 or more). About half of the sample owned stocks with varying degrees of numeracy and financial literacy (the standard deviation of these variables, which corresponds to the fraction of correct answers in numeracy and financial literacy tests, is respectively 0.25 and 0.24).

Table 1: Summary Statistics

	<i>Means by Treatment Arm</i> (<i>Std. deviation in parenthesis</i>)				<i>P-value Test</i>			
	I	II	III	IV	I = III	I = IV	II = IV	III = IV
<u><i>Individual Characteristics</i></u>								
Age*	48.7 (13.74)	47.8 (14.74)	48.4 (16.35)	47.2 (14.71)	0.82	0.30	0.70	0.48
{Male}	0.48	0.45	0.47	0.51	0.96	0.56	0.29	0.54
Numeracy*	0.49 (0.27)	0.46 (0.24)	0.49 (0.24)	0.49 (0.25)	0.90	0.80	0.39	0.91
Financial Literacy*	0.71 (0.24)	0.67 (0.23)	0.68 (0.24)	0.72 (0.23)	0.15	0.87	0.08	0.10
{Own Stocks*}	0.49	0.52	0.49	0.51	0.97	0.75	0.83	0.79
CCEI at Baseline*	0.88 (0.14)	0.90 (0.13)	0.88 (0.16)	0.90 (0.13)	0.97	0.31	0.94	0.39
Risk Aversion at Baseline*	0.67 (0.13)	0.67 (0.13)	0.68 (0.14)	0.68 (0.13)	0.42	0.52	0.75	0.83
<u><i>Education</i></u>								
{Less than High School*}	0.03	0.07	0.05	0.03	0.44	0.73	0.08	0.26
{High School Graduate*}	0.17	0.20	0.16	0.16	0.70	0.69	0.32	1.00
{Some College*}	0.20	0.22	0.24	0.21	0.33	0.80	0.86	0.47
{College Graduate*}	0.60	0.52	0.55	0.61	0.41	0.83	0.09	0.30
<u><i>Annual Household Income</i></u>								
{Less than \$10,000}	0.07	0.07	0.06	0.04	0.75	0.14	0.22	0.27
{Between \$10,000 and \$20,000}	0.10	0.04	0.08	0.07	0.70	0.28	0.39	0.51
{Between \$20,000 and \$30,000}	0.10	0.12	0.12	0.13	0.55	0.49	0.95	0.95
{Between \$30,000 and \$40,000}	0.06	0.12	0.10	0.10	0.11	0.10	0.67	1.00
{Between \$40,000 and \$50,000}	0.09	0.10	0.10	0.06	0.84	0.27	0.15	0.20
{Between \$50,000 and \$60,000}	0.07	0.08	0.06	0.08	0.58	0.78	0.91	0.40
{Between \$60,000 and \$75,000}	0.10	0.16	0.17	0.15	0.08	0.20	0.80	0.59
{Between \$75,000 and \$100,000}	0.18	0.12	0.16	0.15	0.62	0.46	0.51	0.83
{Between \$100,000 and \$150,000}	0.11	0.13	0.07	0.15	0.19	0.34	0.62	0.03
{More than \$150,000}	0.11	0.06	0.07	0.08	0.19	0.31	0.34	0.72
<i>Observations</i>	178	181	158	183				

Notes: This table reports summary statistics and tests whether controls are balanced across the different treatment arms. The first four columns report means for each treatment arm. The standard deviations of continuous variables are reported between parentheses. The last four columns report p-values of tests of the differences in means. Curl brackets indicate dichotomous variables. Asterisks indicate the 10 variables that were used in the re-randomized procedure.

The last four columns of Table 1, which present the p-values of tests of differences in means, show that the observable characteristics are orthogonal to treatment assignment. Out of 84 comparisons, 4 are significant at 10% and one is significant at 5%.

3.2 Portfolio Choices

Table 2 examines the effects of complexity on portfolio choice by comparing the expected return and volatility of the portfolios selected in treatment arm III (complex without outside option) to

those of the portfolios selected in arm I (simple without outside option).¹⁶ It presents results from OLS regressions of the dependent variables listed in the columns – the expected return in dollars, the log of expected return, the rate of return (the expected return as a fraction of the endowment) multiplied by 100, and the standard deviation of the portfolio – on an indicator for being assigned to treatment arm III and a constant. Standard errors are clustered at the individual level.

Table 2: Effects of Complexity on Portfolio Choices

	<i>Expected Return</i>	<i>Ln(Expected Return)</i>	<i>Rate of Return * 100</i>	<i>Standard Deviation</i>
{Complexity}	-\$1.27 [0.40]***	-0.05 [0.02]***	-7.99 [2.32]***	-\$2.08 [0.86]**
Constant	\$28.25 [0.29]***	3.28 [0.01]***	19.75 [1.69]***	\$12.09 [0.64]***

Notes: This table compares the portfolio choices in treatment arm III (complex without outside option) to the portfolio choices in treatment arm I (simple without outside option). Curly brackets indicate dichotomous variables. Standard errors clustered at the individual level are in brackets. The analysis excludes 275 choice sets where all portfolios yield the same expected return. N Choices = 8,125. N Participants = 336.

Complexity leads participants to choose portfolios with lower return and lower risk. The portfolios selected in the complex condition have an expected return that is \$1.27 lower, on average, than the portfolios selected in the simple condition, corresponding to a 4%-5% decrease. On average, portfolios selected in the complex condition have a rate of return 8 percentage points lower than those in the simple condition. All of these differences are statistically significant at the 1% confidence level. The standard deviation of the portfolios selected in treatment arm III is, on average, \$2.08 lower than of the portfolios selected in treatment arm I.

To put these estimates into perspective, Online Appendix Table 1 estimates the cross-sectional relationship between having a college degree and portfolio choices (the sample is restricted to treatment arm I – simple without outside option). The effect of complexity corresponds approximately to one-half of the “returns to a college degree.”

3.3 Decision-making Quality

The preceding results show that participants tend to take less risk and earn lower returns in the complex setting. The theory of “financial competence,” introduced by Ambuehl, Bernheim, and

¹⁶Participants assigned to the complex problem spend substantially more time making decisions. Comparing treatment arms I and III, we find that participants assigned to the simple condition spent on average 10 minutes and 40 seconds making choices, and participants assigned to the complex condition spent on average 19 minutes and 56 seconds. This difference in mean time spent is statistically significant at the 1% level.

Lusardi (2014), interprets these effects as the result of lower-quality decision-making, and the larger the effect of complexity the lower the financial competence. The Ambuehl, Bernheim, and Lusardi (2014) approach compares the choices an individual makes when a decision problem is framed simply to his choice when the same decision problem is framed in a complex manner. Choices in the simple frame are benchmarks; the larger the gap between simple- and complex-framed choices, the lower the individual’s financial competence.

We take a different approach. As described in section 2.5, we measure decision-making quality by compliance with normative properties of choice and demand (transitivity, monotonicity, symmetry). To evaluate complexity’s influence on decision-making quality, Table 3 compares 4 measures of the quality of the choices made in treatment arm III (complex without outside option) to those of the choices made in treatment arm I (simple without outside option). See section 2.5 for how these measures are constructed. With the exception of the propensity to make dominated choices (third column), higher values correspond to higher-quality decisions.

Table 3: Effects of Complexity on Decision-Making Quality

	GARP <i>CCEI</i>	GARP+FOSD <i>CCEI</i>	% <i>Dominated</i> <i>Portfolio</i>	FOSD <i>FOSD Score</i>
{Complexity}	0.03 [0.02]	-0.03 [0.03]	0.09 [0.02]***	-0.01 [0.01]**
Constant	0.86 [0.02]***	0.69 [0.02]***	0.28 [0.02]***	0.94 [0.01]***
P-value Wilcoxon	0.62	0.02	0.00	0.00

Notes: This table investigates if complexity affects the quality of decision-making. It compares measures of the decision-making quality of treatment arm III (complex without outside option) to the decision-making quality of treatment arm I (simple without outside option). Curly brackets indicate dichotomous variables. Robust standard errors are in brackets. The last two columns exclude choice sets where all portfolios yielded the same expected return. N Participants = 336.

There is little evidence that complexity induces more violations of transitivity. Participants in treatment arm III comply a bit more closely with GARP than treatment arm I, but this difference is not statistically significant.¹⁷

Compliance with GARP may be viewed as a necessary but not sufficient condition for high-quality decision-making. Analysis of a unified measure of violations of GARP, FOSD and symmetry of demand via an evaluation of both the actual choices and their mirror image (second column), shows that complexity modestly reduces this measure of decision-making quality. The difference

¹⁷Online Appendix Figure 4 shows the cumulative distribution of the CCEI score, separately for treatment arms I and III. It illustrates that this result is mostly driven by a difference in mass at lower levels of CCEI.

in means is not statistically significant, but we can reject the null of a Wilcoxon rank-sum test at 5%. That is, participants assigned to the complex condition have on average lower ranks (i.e., lower decision-making quality) in the distribution of the unified measure of violations of GARP, FOSD and symmetry of demand, than participants assigned to the simple condition.¹⁸

Isolated analysis of violations of monotonicity with respect to first-order stochastic dominance shows more definitively the effects of complexity on decision-making quality. The third column of Table 3 shows that participants assigned to the complex condition are 9 percentage points more likely to pick a dominated portfolio than participants assigned to the simple condition. The difference in means in the FOSD score (last column), which is statistically significant at 5%, confirms this result. To put into perspective, Online Appendix Table 2 shows that participants with a college degree are 14 percentage points less likely to pick a dominated portfolio than their peers.¹⁹

3.4 The Option to Avoid Complex Problems

The preceding analysis shows that complexity has negative effects on the quality of decision-making, especially on the tendency to pick dominated portfolios. These pitfalls of complexity might be avoided, however, if individuals are sophisticated and know when they should choose simple alternatives rather than solve complex problems. To evaluate the consequences of opting out, here we study the effects on expected returns and decision-making quality of giving participants the option to avoid complex portfolio problems.

Specifically, in Table 4 we study the effects of offering the option to avoid complex portfolio problems on the expected payoff, the log of the expected payoff, the rate of return (i.e., the expected payoff as a fraction of the endowment) multiplied by 100, and compliance with FOSD (measured by the FOSD score). If a participant chose to invest, the expected payoff is equal to the expected return and the FOSD score is as defined above (Section 2.5). If a participant chose to avoid, the expected payoff is equal to the outside option and the FOSD score is equal to $\min\{1, (\text{outside option})/(\text{risk-free return})\}$. The coefficients on the outside option indicator in Table 4 compare the outcomes in treatment arm IV (complex with outside option) to those in treatment arm III (complex without outside option). The constant terms simply describe the average

¹⁸Angrist and Imbens (2009) argue that “[i]f the focus is on establishing whether the treatment has some effect on the outcomes, rather than on estimating the average size of the effect, such rank tests [as the Wilcoxon] are much more likely to provide informative conclusions than standard Wald tests based differences in averages by treatment status. . . As a general matter it would be useful in randomized experiments to include such results for rank-based p-values, as a generally applicable way of establishing whether the treatment has any effect.” (pp. 22-23)

¹⁹Some of the degradation in decision-making quality might be attributed to additional reliance on heuristics or rules of thumb for making decisions. An especially important rule is the “1/n strategy,” in which individuals divide their wealth equally among n investment options (Benartzi and Thaler, 2001). Because participants choose a number of shares rather than a direct dollar allocation, this heuristic may not be as compelling in this setting. Regardless, we find no evidence that complexity led to increases in the tendency to equate either the amount of money invested in each asset or to equate the number of shares purchased of each asset. Both of these behaviors are rare, representing less than 5% of choices.

levels of these outcomes in treatment arm III and thus summarize the results of Table 2.

Table 4: Effects of the Option to Avoid Complex Decision-Making

	<i>Expected Payoff</i>	<i>Ln(Expected Payoff)</i>	<i>Rate of Return * 100</i>	<i>FOSD Score</i>
{Outside Option}	-\$2.21 [0.47]***	-0.15 [0.03]***	-8.99 [2.48]***	-0.06 [0.01]***
Constant	\$26.97 [0.27]***	3.22 [0.01]***	11.76 [1.59]***	0.93 [0.01]***

Notes: This table investigates if having the option to avoid complexity mitigates its effects. It compares the payoffs of treatment arm IV (complex with outside option) to the payoffs of treatment arms III (complex without outside option). The payoff is equal to the outside option if the participant chose to avoid the investment decision-making and equal to the portfolio return if the participant chose to invest. Curly brackets indicate dichotomous variables. For participants in treatment arm IV who chose to avoid complexity, the FOSD score is equal to the outside option divided by the return of the risk-free portfolio if outside option < return of risk-free portfolio and equal to 1 otherwise. Standard errors clustered at the individual level are in brackets. N Choices = 8,228. N Participants = 341. We exclude choice sets where all portfolios yielded the same expected return.

We find no evidence that the ability to opt out helps participants avoid worse choices in the complex portfolio problem. Instead, the availability of the outside option amplifies the negative effects of complexity. It lowers the rate of return by 9 percentage points (relative to those in the complex condition with no outside option). Expected payoffs reduce by 15 percent, three times as much as the effect of complexity when there is no outside option (see Table 2). A reduction in the quality of decision-making, in particular in the compliance with monotonicity with respect to FOSD, underlies these effects.

The ability to opt out may hurt those with lower decision-making skills the most if they do not know when they should choose the outside option. We study the effects of the simple alternative on the lower-skilled by, first, calculating a predetermined measure of decision-making skills. Specifically, we identify decision-making skills with the first component of a principal component analysis of three variables: the score in a numeracy test, the score in a financial literacy test, and consistency with GARP measured at baseline. We stratified randomization of the experimental treatments on these three variables in anticipation of analysis examining whether the effects of complexity vary by decision-making skills. The first principal component was re-scaled to range from 0 to 1. Second, we investigate how the consequences of the ability to opt out vary by decision-making skill. Table 5 compares the choices of treatment arm IV (complexity with outside option) to the choices of treatment arm III (complexity without outside option), allowing the effect of the outside option to vary with decision-making skills.

When offered the outside option, participants with the lowest decision-making skills have an expected payoff 40 percent lower than they would have otherwise. The penalty associated with the

option to avoid complexity is especially large for those with the least decision-making skills because there is a large reduction in their compliance with a dominance principle. In contrast, those with high decision-making skills guard themselves against the negative effects of having the outside option. The coefficients on the interaction terms are positive and the point estimates indicate that the effects of the outside option for someone with the highest level of decision-making skills are close to zero.

Table 5: Effects of the Option to Avoid, by Decision-making Skill (Complex Condition)

	<i>Expected Payoff</i>	<i>Ln(Expected Payoff)</i>	<i>Rate of Return * 100</i>	<i>FOSD Score</i>
Decision-making Skill * {Outside Option}	\$4.12 [2.25]*	0.38 [0.16]**	24.33 [12.18]**	0.13 [0.06]**
{Outside Option}	-\$4.86 [1.41]***	-0.39 [0.11]***	-24.47 [7.19]***	-0.15 [0.04]***
Decision-making Skill	\$4.88 [1.27]***	0.21 [0.05]***	24.23 [6.68]***	0.12 [0.03]***
Constant	\$24.11 [0.69]***	3.10 [0.03]***	-2.45 [3.48]	0.86 [0.02]***

Notes: This table investigates if the effects of having the option to avoid complexity differ by decision-making skills. It compares the payoffs of treatment arm IV (complex with outside option) to the payoffs of treatment arm III (complex without outside option). The payoff is equal to the outside option if the participant chose to avoid the investment decision-making and equal to the portfolio return if the participant chose to invest. Curly brackets indicate dichotomous variables. For participants in treatment arm IV who chose to avoid complexity, the FOSD score is equal to the outside option divided by the return of the risk-free portfolio if outside option < return of risk-free portfolio and equal to 1 otherwise. Standard errors clustered at the individual level are in brackets. N Choices = 8,203. N Participants = 340. We exclude choice sets where all portfolios yielded the same expected return and dropped 1 participant for whom numeracy or financial literacy scores were missing.

Table 6 investigates why the least-skilled comply less with a dominance principle when they can opt out. Is it because they opt out even when the outside option pays less than the risk-free portfolio? Or does the option to opt out lead them to make worse investments which are dominated (by some other feasible portfolio)? The table compares the likelihood in treatment arms IV (complexity with outside option) and III (complexity without outside option) of making a dominated choice, allowing the effect of the outside option to vary with decision-making skills. The first column looks at the likelihood of making any dominated choice, either by picking a dominated portfolio or by opting out when when the outside option is dominated by the risk-free portfolio. The last two columns distinguish between these two types of dominated choices.

When the outside option is made available, there is a similar increase in the likelihood that the least- and highest-skilled make dominated choices. The least- and highest-skilled are 8 and 9 percentage points more likely to violate monotonicity with respect to FOSD. But these results mask important heterogeneity. The ability to opt out increases the chance the highest-skilled make a dominated investment and the chance they opt out when the outside option is dominated by similar amounts: 4 and 5 percentage points. In contrast, for the least-skilled, there is a large

shift from dominated portfolios to dominated outside options. They are 30 percentage points less likely to make a dominated investment and 38 percentage points more likely to make a dominated avoidance decision. When the simple alternative is introduced, the least-skilled are only 5 percentage points more likely than the highest-skilled to make a dominated investment. They are, however, 33 percentage points more likely to opt out when doing so is dominated.

Table 6: Dominated Choices and the Option to Avoid

	Dominated Choice	Dominated Portfolio	Dominated Avoidance
Decision-making Skill * {Outside Option}	0.01 [0.09]	0.34 [0.09]***	-0.33 [0.08]***
{Outside Option}	0.08 [0.06]	-0.30 [0.06]***	0.38 [0.06]***
Decision-making Skill	-0.39 [0.07]***	-0.39 [0.07]***	0.00 -
Constant	0.60 [0.04]***	0.60 [0.04]***	0.00 -

Notes: This table investigates if the least-skilled comply less with a dominance principle because they opt out when the outside option is dominated or because they pick dominated portfolios. It compares the likelihood of treatment arm IV (complex with outside option) and treatment arm III (complex without outside option) of making dominated choices. N Choices = 8,203. N Participants = 340. We exclude choice sets where all portfolios yielded the same expected return.


4 Sophistication

In a complex setting, the availability of a simple alternative to solving a portfolio problem results in sharply lower returns and more dominated choices for the lower-skilled. These negative consequences of the simple alternative derive from the fact that these participants opt out more often, even when doing so is dominated. In one view, these results imply a lack of sophistication; the lower-skilled appear not to know when they are better off taking a simple alternative to solving a complex problem. This interpretation of the evidence does not, however, account for costs of attending to the portfolio problem. It is plausible that lower-skilled participants face higher costs of obtaining information about the varying, payoff-relevant elements of the portfolio problem. If so, then they might rationally opt out and trade attention costs for lower returns.

Metadata on participants' choices are consistent with lower-skilled participants facing higher costs of attending to complex problems. Even when they have no opportunity to opt out of complexity, those with lower skills make fewer active adjustments to their portfolios and spend substantially less time making their choices. To see this, for each of a participant's 25 choices we

calculated the number of times that the participant changed either from increasing to decreasing or from decreasing to increasing the amount received if the coin came up heads. We then calculated the median of this number across the participant's 25 decisions. Splitting the sample in half by decision-making skill, the top panel of Figure 2 shows the cumulative distribution of these medians across participants. Among those with lower skills, 60% had a median number of adjustments less than 2. This number is about 35% among those with higher skills. Similarly, we measured the amount of time a participant spent on each problem and calculated the median of this number across the participant's 25 decisions. The bottom panel of Figure 2 shows the cumulative distribution across participants of that median, separately for the low- and high-skilled. Among low-skilled participants, about 80% had a median time of less than a minute. The equivalent number is 57% for the high-skilled. Making fewer active adjustments and spending less time, lower-skilled participants appear to pay less attention to the problems, perhaps because they face higher costs of doing so.

Figure 2: Distributions of Interface Use and Time Spent, By Skill



Notes: The top panel shows cumulative distribution functions of interface use, separately for the low-skilled ($N = 77$) and the high-skilled ($N = 81$). For each one of a participant's 25 choices, we first calculated the number of times that the participant changed from increasing to decreasing the amount received if the coin came up heads or changed from decreasing to increasing the amount received if heads and then calculated the median across the

participant’s 25 decisions. The CDFs shows the distributions of these medians across participants. The bottom panel shows cumulative distribution functions of the time participants spent choosing their allocations, separately for the low- and high-skilled. We first calculated the median amount of minutes a participant spent choosing across her 25 choices. The CDFs show the distributions of these medians across participants. The sample is restricted to participants assigned to treatment arm III (complex without outside option).

4.1 Rational Inattention Model

The metadata described above suggest that higher attention costs may account for the higher rates of opting out by those with lower skills. To draw further inference about the differences in attention costs by skill, and to evaluate further the hypothesis of sophisticated opting out, we estimate structural parameters of a rational inattention model based on Sims (2003) and developed in Matějka and McKay (2015). For this purpose, we develop a novel and portable method for identification that allows sufficiently precise estimation of attention cost parameters with relatively mild assumptions and data requirements.

In this model, the participant is uncertain about the value of each of the options she may choose, and has a prior belief about those values. The participant adopts an optimal information acquisition (attention) strategy by which she obtains knowledge about the values of her options and updates her prior. Attention is costly. Inference does not require that the attention strategy be specified. A strategy might, for example, describe which aspects of the problem a participant attends to and to what extent. The strategy could specify whether and how carefully to attend to the number of assets in the choice set, the relative prices of assets, the payoffs of each asset, the endowment she has to spend, the outside option level, etc. Regardless of strategy, the model assumes that, based on her posterior belief about the value of her options, the participant chooses the one with the highest expected utility.

Formally, we follow the structure and notation of Matějka and McKay (2015). We restrict attention to treatment arm IV (complex with outside option) and model each decision problem as presenting the participant with N options indexed by i . Option $i = N$ is to opt out and take the simple alternative. The remaining $N - 1$ options are the elements of the set of feasible portfolios. As described below, the empirical strategy we implement requires only limited assumptions about how that set of feasible portfolios is specified. The value to the participant of each option i , denoted v_i , is uncertain. Let $\mathbf{v} = (v_1, \dots, v_N) \in \mathbb{R}^N$ denote this uncertain state. We assume the participant is endowed with a prior belief about the distribution of \mathbf{v} , $G \in \Delta(\mathbb{R}^N)$, where $\Delta(\mathbb{R}^N)$ is the set of all probability distributions on \mathbb{R}^N .

The participant knows her information about the state is imperfect. Before choosing an option, therefore, the participant selects an attention strategy to refine her prior. We assume the strategy reduces uncertainty about the state (\mathbf{v}) and results in a posterior belief $F \in \Delta(\mathbb{R}^N)$. The uncer-

tainty of beliefs is described in terms of entropy. If $H(G)$ is the entropy of the prior G , if the state distribution is discrete, and if p_k is the probability of state k , then $H(G)$ satisfies

$$H(G) = - \sum_k p_k \log(p_k).$$

Entropy thus gives the average log likelihood of each state. If, for example, there were just two states then entropy would rise with variance and is maximized when each state is equally likely.

Attention is costly. Following Sims (2003) and the literature that follows, we assume the costs of attention are linear in entropy reduction. Starting with prior G and associated uncertainty $H(G)$, to arrive at posterior beliefs F with associated uncertainty $H(F)$ involves a cost $c(F)$ that satisfies

$$c(F) = \lambda [H(G) - H(F)].$$

We assume the participant chooses an attention strategy to maximize the expected utility of her choice net of the costs of that strategy $c(F)$.

Matějka and McKay (2015) show that optimal behavior in this model implies the conditional probability a participant chooses option i , $P(i, \mathbf{v})$, satisfies

$$P(i, \mathbf{v}) = \frac{e^{(v_i + \alpha_i)/\lambda}}{\sum_{j=1}^N e^{(v_j + \alpha_j)/\lambda}} \quad (1)$$

where α_i is the prior weight assigned to option i .²⁰ ²¹ The prior weight α_i describes the relative tendency to choose option i in the absence of additional information about its actual value v_i .

The optimal conditional choice probabilities in (1) imply that when attention costs λ are high, the prior weights α_i dominate the true values v_i . This is perhaps easiest to see in the ratio of the conditional probabilities of choosing any two options i and j . This ratio is given by

$$\frac{P(i, \mathbf{v})}{P(j, \mathbf{v})} = \frac{e^{(v_i + \alpha_i)/\lambda}}{e^{(v_j + \alpha_j)/\lambda}} = \frac{P(i)}{P(j)} * e^{(v_i - v_j)/\lambda} \quad (2)$$

where the unconditional probability of choosing option i , $P(i) \equiv \int_{\mathbf{v}} P(i, \mathbf{v}) G(d\mathbf{v}) = e^{\alpha_i/\lambda}$.

Equation 2 shows that when the costs of attention are very high, the difference in the realized values of different options, $(v_i - v_j)$, has little influence on the relative probabilities of choosing those options. Instead, the relative likelihood of choosing option i instead of j is driven by the

²⁰Caplin, Dean, and Leahy (2019) emphasize that this conditional choice probability formula holds only for options the individual believes worth considering. It holds, that is, only for those options with a positive unconditional choice probability. Our estimation strategy is sufficiently flexible to allow some options to have zero unconditional choice probabilities. We require only that the unconditional probability of opting out is not zero or one.

²¹Matějka and McKay (AER 2015) studies which elements of the rational inattention model distinguish it from a standard, multinomial logit model of random choice from a finite number of options.

unconditional probabilities. Choices are driven, that is, by the average probabilities of making either choice before any information is acquired. Conversely, when attention costs are low, true values dominate and as λ gets arbitrarily small the probability of choosing the option with the highest true value goes to 1.

The model thus offers an interpretation of differences in rates of opting out in the experiment. Specifically, and other things equal, if lower-skilled participants are less responsive to the true value of each option, this would be captured by a higher cost of attention λ . At the same time, it is sufficiently flexible to accommodate alternative explanations. For example, it can incorporate salience *a la* Bordalo, Gennaioli, and Schleifer (2012). If the outside option stood out from the other options in terms of some non-pecuniary attribute (e.g., its position on the screen), prior weights for options that fare well (poorly) in the terms of the salient attribute would increase (decrease) to reflect the distortion of decision weights in favor of the salient attribute. The prior weights on the outside option could also capture any differential beliefs between the lower- and higher-skilled (e.g., if the lower-skilled were under the impression that the experimenter endorsed the outside option).

4.2 Identification, Estimation, and Interpretation

It is difficult to identify separately v_i , α_i , and λ from choice data alone. Separate identification is challenging because none of these parameters is observed directly and, given a cost of attention (λ), the realized value of option i (v_i) and the prior weight on that option (α_i) have similar influence on the probability of choosing that option. In the controlled experimental setting, however, we can make progress by exploiting random variation in the contingent payoffs that determine the fundamental value of an option, v_i .

The logic for identification is, again, evident in the ratio of the conditional probabilities of choosing any two options i and j :

$$\frac{P(i, \mathbf{v})}{P(j, \mathbf{v})} = \frac{P(i)}{P(j)} * e^{(v_i - v_j)/\lambda} \quad (3)$$

where the unconditional probability of choosing option i , $P(i) = e^{\alpha_i/\lambda}$. Taking logs of both sides of equation (3) we obtain

$$\ln \left(\frac{P(i, \mathbf{v})}{P(j, \mathbf{v})} \right) = \ln \left(\frac{P(i)}{P(j)} \right) + \frac{v_i - v_j}{\lambda} \quad (4)$$

As equation (4) makes clear, the change in the log likelihood ratio of choosing option i versus j caused by a change in the difference in their values is given by the inverse of λ . Thus, random variation in the contingent payoffs that determine v_i in v_1, \dots, v_N can be used to identify the cost of attention. If λ is thus identified, then the prior weights α_i could be backed out from the average

log likelihood ratios.²²

This method of identification relies, however, on several strong assumptions and imposes substantial data demands. These assumptions and data demands include:

1. Specification of the set of available options i in each problem. In our context, this means an assumption about how the space of available portfolios is partitioned.
2. Assumptions on the functional form of the utility function that maps the contingent payoffs of each option i into its value v_i .
3. Sufficiently large samples to estimate $N - 1$ parameters: $\alpha_1, \dots, \alpha_{N-1}$.

We avoid these assumptions and data demands by exploiting the fact that, if preferences are homothetic, v_1, \dots, v_{N-1} are proportional to the dollar amount of the endowment while v_N is proportional to the dollar amount of the outside option. This assumption allows identification of λ without further functional form assumptions on the utility function.

To see how, consider for the moment a setting where asset prices are fixed, the experimental endowment is m and the outside option is ω . The identification method requires that each of the $N - 1$ options that represent the set of feasible portfolios expands proportionally with m . Formally, if $\mathbf{x}_i(1)$ denotes option i when $m = 1$, then $\mathbf{x}_i(m) = m\mathbf{x}_i(1)$. For example, the option \mathbf{x}_i may describe the amount invested in each asset or the numbers of shares of each asset. We also assume individuals view opting out of the portfolio problem, option N , as one of the available options.

Under these assumptions, the ratio of the conditional probability of opting in to the conditional probability of opting out can be written as:

$$\frac{1 - P(N, \mathbf{v}(m, \omega))}{P(N, \mathbf{v}(m, \omega))} = \sum_{j=1}^{N-1} \frac{P(j)}{P(N)} * e^{\{\ln v_j(m) - \ln v_N(\omega)\} / \lambda'} \quad (5)$$

where $\ln v_j$ is a log transformation of v_j and λ' is the appropriate rescaling of attention costs to reflect the cardinality of $\ln \mathbf{v}$.

If preferences are homothetic, $\ln v_j(m) = \ln m + \ln v_j(1)$ and $\ln v_N(\omega) = \ln \omega + \ln v_N(1)$, such that the equation above can be rewritten as:

$$\frac{1 - P(N, \mathbf{v}(m, \omega))}{P(N, \mathbf{v}(m, \omega))} = e^{\{\ln m - \ln \omega\} / \lambda'} \left[\sum_{j=1}^{N-1} \frac{P(j)}{P(N)} * e^{\{\ln v_j(1) - \ln v_N(1)\} / \lambda'} \right] \quad (6)$$

²²An alternative approach would elicit the unconditional probability of choices $P(i)$ directly from participants and use this as data with which to estimate the other parameters of the model. Note that $P(i)$ is not, however, simply the prior belief about the optimality of choosing option i . It is, instead, the prior on choosing option i accounting for both the uncertainty of the state and the optimal attention strategy. The identification strategy described below avoids the challenge and expense of having to elicit this seemingly abstract object.

Notice that the term between brackets is constant (for a given set of asset prices), which can be represented as a fixed effect in a conditional logit.²³

It follows that λ' can be identified from the sensitivity of the ratio of the conditional probabilities to the random variation in the endowment and in the outside option. Intuitively, the identification strategy approximates an experiment where treatment and control, who face the same asset prices and the same returns, are assigned different endowments or outside options. Based on this identification strategy, we can estimate the parameters via maximum likelihood and allow the parameters to vary with decision-making skill. Notice that the formulation in equation (6) allows the lower- and higher-skilled to have different time and risk preferences, which would be captured by the skill-specific fixed effects.

For purposes of evaluating sophistication we will interpret the structural estimates as follows. In the model, there are three reasons why lower-skilled participants may opt out more often than those with higher skills, even when opting out is dominated. It is either because the low-skilled derive a different relative value from opting out (v_N), or because they place a different relative prior weight on opting out (α_N), or because they have higher costs of attention (λ). Any of these reasons for seemingly dominated opting out is, by construction, rational. We will classify the higher rates of seemingly dominated opting out by the lower-skilled as sophisticated only to the extent they are driven by higher costs of attention. In this view, sophistication means opting out at a monetary cost in order to save the utility costs of becoming more certain about which is the best option.

4.3 Results

Table 7 presents estimates of the cost of attention (λ), by skill, in complex problems. Here high- and low-skilled are defined by being above or below the median level of the decision-making skill index. The first column shows the point estimate of λ for low-skilled participants (second row) and the point estimate of the difference in λ by skill (first row). The last two columns show the bounds of 95% confidence intervals for these estimates, calculated by bootstrapping. The point estimates indicate that lower-skilled participants have a cost of attention that is more than twice as large as that of higher-skilled participants. The confidence intervals around these point estimates indicate, as expected, that we can reject a null hypothesis of costless attention ($\lambda = 0$), and a null hypothesis of equal costs of attention between skill groups.

²³ In the estimation, we allow for fixed effects specific to each set of asset prices \times decision-making skill combination.

Table 7: Structural Estimates of Attention Costs, by Skill (Complex Problems)

	Point Estimate	Bounds of bootstrapped 95% confidence interval	
		<i>Lower</i>	<i>Upper</i>
High Skilled - Low Skilled	-0.900	-2.957	-0.017
Low Skilled	1.660	0.888	3.644

Notes: This table presents maximum likelihood estimates of the costs of information acquisition parameter, λ , for low-skilled participants and the difference in this parameter between high- and low-skilled participants. Attention is restricted to complex treatments in which the outside option is available. N Choices = 8,203. N Participants = 340.

We interpret the results in Table 7 as consistent with the hypothesis that the higher rate of costly opting out among lower-skilled participants in complex problems is, to some extent, sophisticated. Viewed through a rational inattention model, the estimates in Table 7 indicate that, relative to higher-skilled participants, the higher rate of dominated avoidance by lower-skilled participants is driven at least in part by their higher costs of attending to their choices.

4.4 Quantifying The Importance of Attention Costs for Compliance with a Dominance Principle

The point estimates in Table 7 show statistically significant differences in the costs of attention between participants of different skills. Because these differences in λ are measured in terms of utility, however, it is not clear that they are economically important. Even large level differences in λ may not translate into large differences in the likelihood of dominated avoidance.

To quantify the ability of the estimated differences in attention costs to explain differential rates of making dominated choices, we first calculate a money metric of the consequences of sub-optimal decision-making – “the money left on the table.” For a given assumption about the utility function, and for each decision problem that a participant faced, the money left on the table is given by the difference between the certainty equivalent of the optimal option and the certainty equivalent of the option the participant actually chose. For each of five different assumptions about the utility function, this calculation is done for each participant and each decision, and then averaged over all their decisions, by skill group. The difference in these averages is a measure of the differences, by skill group, in the money being left on the table by sub-optimal choices. The first row of Table 8 presents these calculations. To illustrate, if we assume risk neutrality, then on average the low-skilled left \$2.85 more on the table than the high-skilled.

Table 8: Amount Left on the Table Explained by Attention Costs, by Skill

	<i>Linear Utility</i>	<i>Leontieff</i>	<i>Log Utility</i>	<i>CES rho = 2</i>	<i>CES rho = 3</i>
High-Skilled - Low-Skilled					
CE Optimal – CE Observed	-\$2.85	-\$1.70	-\$1.58	-\$1.54	-\$1.65
CE Optimal – CE Simulated	-\$1.10	-\$2.29	-\$0.58	-\$1.78	-\$2.06
Low-Skilled					
CE Optimal – CE Observed	\$14.71	\$9.03	\$6.47	\$7.21	\$7.71
CE Optimal – CE Simulated	\$8.61	\$9.63	\$5.81	\$6.81	\$7.68

Notes: This table evaluates how well the between-skill-group difference in money left on the table can be explained by the between-skill-group difference in estimated costs of information acquisition (λ). It first calculates, for five different assumptions about the functional form of utility (assumed to be the same for both groups), the between-skill-group difference in the observed amount of money left on table (first row). Then it assumes uniform unconditional probabilities and simulates, for each of the utility functions, the amount of money left on the table predicted by the model with the estimated λ s (second row). The third and fourth rows present the levels of the observed and simulated for the low skilled group.

To what extent can the estimated differences in attention costs (λ) explain these skill group differences in the money left on the table? To answer this question, in addition to a functional form for utility, we need to impose more assumptions we avoided in estimation. First we need to assume a specific partition of the option space inside the portfolio. Here we assume the option space for each problem is partitioned according to the fraction of the endowment invested in each Arrow security. Second, we need to accommodate the influence of attention costs on both the conditional probability of choosing option i , $P(i, \mathbf{v})$, and the unconditional probability of choosing option i , $\int_{\mathbf{v}} P(i, \mathbf{v})G(d\mathbf{v})$. To our knowledge, the latter, which depends on prior beliefs $G(\mathbf{v})$, has no closed form. For the calculations reported below, therefore, we assume that, despite their different attention costs, both the high- and low-skilled have uniform unconditional probabilities.²⁴ The calculations thus isolate the ex-post effect of attention costs on conditional choice probabilities, separate from their influence on the probability of making a choice before any information is acquired.

The assumption of uniform unconditional probabilities yields:

$$P(i, v) = \frac{e^{\ln v_i / \lambda'}}{\sum_{j=1}^N e^{\ln v_j / \lambda'}} \quad (7)$$

Given a functional form of the utility function, we use equation (7) and the estimated λ s to simulate the choice behavior of participants. The simulated amount left on the table is calculated as the

²⁴ For them to share the same unconditional probabilities, and yet have different costs of attention, the low- and high-skilled must have different prior beliefs $G(\mathbf{v})$.

difference between the certainty equivalence of the optimal option and the certainty equivalence of the simulated choice.

The second row of Table 8 shows the differences in attention costs explain a substantial fraction of the difference in the amounts foregone by low- and high-skilled participants. If we assume risk neutrality for example, the differences in λ s predict the low-skilled will leave (on average) \$1.10 more on the table than the high-skilled, explaining 39% of the extra amount actually foregone by the low-skilled. Depending on the utility function, this percentage explained ranges from 37% for log utility to 135% for Leontieff preferences. The estimated attention costs also match well the levels of money left on the table as evidenced by the comparison of the last two rows of the table.

These results show that the estimated differences in costs of attention are sufficiently large to explain the differences by decision-making skill in compliance with a dominance principle. In this way, they suggest that the low-skilled may often be sophisticated in their decision to take a simple alternative to solving a complex problem. The estimated model implies that they may forego this money because doing so saves the costly effort of attending to the problem more thoroughly.

5 Conclusion

Evolving financial products and investment opportunities can provide more people greater autonomy and access to the benefits of financial markets. This potential may be limited, however, if consumers are poorly equipped to handle the increased complexity associated with the new choices. Providing such consumers with simple alternatives, like target-date retirement saving plans, or age-based college savings plans, is a sensible way to guard against some negative effects of increasingly complex financial markets. The benefits of these simple alternatives may depend, however, on consumer sophistication. If they can now avoid complex financial decisions, it becomes important for consumers to know when they are better off choosing simple options instead of solving complex problems. Are consumers sufficiently self-aware to see when they ought to avoid complexity in favor of a simple, perhaps imperfect, alternative?

This paper describes an experiment, conducted with a large and diverse population of Americans, that evaluates the effects of complexity on financial choices and assesses the sophistication of individuals to know when they are better off taking a simple option instead of solving a complex problem. Consistent with concerns about the influence of complexity, the results show that, when they are required to make an active portfolio decision, participants facing complex problems make choices with lower expected payoffs and lower risk. On average, complexity also reduces some desirable properties of choice; it leads especially to more violations of monotonicity with respect to first-order stochastic dominance.

The availability of the simple alternative amplifies the negative effects of complexity rather than help participants avoid worse choices. Those with the lowest levels of skills are hurt the most,

earning much lower returns because they often opt out when the outside option is dominated. They appear therefore unsophisticated, not knowing when they are better off opting out. If, however, lower-skilled participants face higher costs of attending to the portfolio problem they may rationally trade these costs for lower returns.

Metadata on choice in the experiment is consistent with higher attention costs for lower-skilled participants. They make fewer active decisions and spend less time on each problem. Attention costs cannot, however, be directly observed. We therefore draw inference about their importance and about sophistication by estimating the structural parameters of a rational inattention model based on Sims (2003) and developed by Matějka and McKay (2015). In this model, a participant is uncertain about the value of each option he faces, but has a prior belief about those values. The participant accumulates costly knowledge about those values and updates her prior. Based on her posterior belief, the participant chooses the option with the highest expected utility.

In our interpretation, if the lower-skilled experience higher costs of attention, we treat the resulting increase in seemingly dominated opting out as sophisticated. In this view, participants are making optimal decisions to opt out at considerable monetary cost rather than incur the higher utility costs of learning more about what, fundamentally, is the best option.

The structural estimates, obtained with a novel and robust identification strategy, are consistent with sophistication. We find a positive, statistically significant, and economically substantial difference between the attention costs of low-skilled and high-skilled participants. The *ex – post* effects of the estimated differences in these attention costs, alone, can explain large fractions of the differences between these groups in the amounts of money they leave on the table by making dominated choices.

Future work should evaluate the robustness of these results in other settings. In the interim, the results of this experiment underscore the importance of taking selection into account when designing simple alternatives to solving complex problems. If lower-skilled people find attending to complex problems too costly, they are more likely to take simple options regardless of their fundamental value. Plan designers should therefore take special care to ensure the simple alternatives are well-suited to the least-skilled who are most likely to take them.

References

- [1] Abeler, Johannes, and Simon Jager. “Complex Tax Incentives.” *American Economic Journal: Economic Policy* 7.3 (2015): 1-28.
- [2] Afriat, Sydney. “Efficiency Estimates of Production Functions.” *International Economic Review* 8 (1972): 568-598.
- [3] Agnew, Julie R., and Lisa R. Szykman. “Asset Allocation and Information Overload: The

- Influence of Information Display, Asset Choice, and Investor Experience.” *Journal of Behavioral Finance* 6.2 (2005): 57-70.
- [4] Al-Najjar, Nabil I., Ramon Casedus-Masanell, and Emre Ozdenoren. “Probabilistic Representation of Complexity.” *Journal of Economic Theory* 111.1 (2003): 49-87.
- [5] Besedeš, Tibor, Cary Deck, Sudipta Sarangi, and Mikhael Shor. “Age Effects and Heuristics in Decision Making.” *Review of Economics and Statistics* 94.2 (2012): 580-595.
- [6] Besedeš, Tibor, Cary Deck, Sudipta Sarangi, and Mikhael Shor. “Decision Making Strategies and Performance among Seniors.” *Journal of Economic Behavior & Organization* 81 (2012): 524-533.
- [7] Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian. “The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States.” In *Social security policy in a Changing Environment*, pp. 167-195. University of Chicago Press, 2009.
- [8] _____. “Simplification and Saving” *Journal of Economic Behavior & Organization*, 95 (2013), 130-145.
- [9] Brocas, Isabelle, Juan D. Carillo, T. Dalton Combs, and Niree Kodaverdian. “Consistency in Simple vs. Complex Choices Over the Life Cycle.” (2014) Working paper, University of Southern California.
- [10] Caplin, Andrew, and Mark Dean. “Revealed Preference, Rational Inattention, and Costly Information Acquisition.” Forthcoming, *American Economic Review* (2015) 105(7):2183-2203.
- [11] Caplin, Andrew, Mark Dean, and Daniel Martin. “Search and Satisficing.” *American Economic Review* 101.7 (2011): 2899-2922.
- [12] Caplin, Andrew, Mark Dean, and John Leahy. “Rational Inattention, Optimal Consideration Sets and Stochastic Choice.” *Review of Economic Studies*, May 2019, 86(3): 1061-1094.
- [13] Carlin, Bruce L., Shimon Kogan, and Richard Lowery. “Trading Complex Assets.” *Journal of Finance* 68.5 (2013): 1937-1960.
- [14] Chetty, Raj, Adam Looney, and Kory Kroft. (2007). “Salience and Taxation: Theory and Evidence.” National Bureau of Economic Research Working Paper 13330.
- [15] Choi, Syngjoo, Raymond Fisman, Douglas Gale, and Shachar Kariv. “Consistency and Heterogeneity of Individual Behavior Under Uncertainty.” *The American Economic Review* 97.5 (2007a): 1921-1938.

- [16] Choi, Syngjoo, Raymond Fisman, Douglas Gale, and Shachar Kariv. “Revealing Preferences Graphically: An Old Method Gets a New Tool Kit.” *The American Economic Review* 97.2 (2007b): 153-158.
- [17] Choi, Syngjoo, Shachar Kariv, Wieland Müller, and Dan Silverman. “Who Is (More) Rational?” *The American Economic Review* 104.6 (2014): 1518-1550.
- [18] Dean, Mark. “Status Quo Bias in Large and Small Choice Sets.” (2008) Working paper, New York University.
- [19] Egan, Mark, Gregor Matvos, and Amit Seru. The Market for Financial Adviser Misconduct. No. w22050. National Bureau of Economic Research (2016).
- [20] Eyster, Erik, and Georg Weizsäcker (2010). “Correlation Neglect in Financial Decision-Making.” Working Paper, London School of Economics and Political Science.
- [21] Friesen, Lana, and Peter E. Earl. “Multipart Tariffs and Bounded Rationality: An Experimental Analysis of Mobile Phone Plan Choices.” *Journal of Economic Behavior & Organization* 116 (2015): 239-253.
- [22] Gale, Douglas, and Hamid Sabourian. “Complexity and Competition.” *Econometrica* 73.3 (2005): 739-769.
- [23] Hadar, Josef, and William R. Russell. “Rules for Ordering Uncertain Prospects.” *American Economic Review* 59.1: (1969): 25-34.
- [24] Huck, Steffen, and Georg Weizsäcker. “Risk, Complexity, and Deviations from Expected-Value Maximization: Results of a Lottery Choice Experiment.” *Journal of Economic Psychology* 20 (1999): 699-715.
- [25] Iyengar, Sheena S., and Mark R. Lepper. “When Choice is Demotivating: Can One Desire Too Much of a Good Thing?” *Journal of Personality and Social Psychology* 79.6 (2000): 995-1006.
- [26] Iyengar, Sheena S., and Emir Kamenica. “Choice Proliferation, Simplicity Seeking, and Asset Allocation.” *Journal of Public Economics* 94.7 (2010): 530-539.
- [27] Iyengar, Sheena S., G. Huberman, and W. Jiang. “How Much Choice is Too Much? Contributions to 401(k) Retirement Plans.” *Pension Design and Structure: New Lessons from Behavioral Finance*. Eds. O. Mitchell and S. Utkus. Oxford: Oxford University Press, 2004. 83-95.
- [28] Kalayci, Kenan, and Marta Serra-Garcia. “Complexity and Biases.” *Experimental Economics* (2012): 1-20.

- [29] Kariv, Shachar, and Dan Silverman. “An Old Measure of Decision-Making Quality Sheds New Light on Paternalism.” *Journal of Institutional and Theoretical Economics* 169.1 (2013): 29-44.
- [30] Kempf, Alexander, and Stefan Ruenzi. “Status Quo Bias and the Number of Alternatives: An Empirical Illustration from the Mutual Fund Industry.” *Journal of Behavioral Finance* 7.4 (2006): 204-213.
- [31] Mador, Galit, Doron Sonsino, and Uri Benzion. “On Complexity and Lotteries’ Evaluation – Three Experimental Observations.” *Journal of Economic Psychology* 21 (2000): 625-637.
- [32] Madrian, Brigitte C., and Dennis F. Shea. “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior.” *The Quarterly Journal of Economics* 116.4 (2001): 1149-1187.
- [33] Matějka, Filip and Alisdair McKay. “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model.” *The American Economic Review* (2015).
- [34] Masatlioglu, Yusufcan, Daisuke Nakajima, and Erkut Y. Ozbay. “Revealed Attention.” *The American Economic Review*, 102.5 (2012): 2183-2205.
- [35] Mullainathan, Sendhil, Markus Noeth, and Antoinette Schoar. The Market for Financial Advice: An Audit Study. No. w17929. National Bureau of Economic Research (2012).
- [36] Ortoleva, Pietro. “The Price of Flexibility: Towards a Theory of Thinking Aversion.” *Journal of Economic Theory* 148(3) (2013): 903-934.
- [37] Phatak, Narahari. “Menu-Based Complexity: Experiments on Choice over Lotteries.” (2012) Unpublished manuscript, University of California at Berkeley.
- [38] Rees-Jones, Alex, and Dmitry Taubinsky (2016). “Heuristic Perceptions of the Income Tax: Evidence and Implications for Debiasing.” NBER Working Paper No. 22884.
- [39] Ren, Yejing. “Status Quo Bias and Choice Overload: An Experimental Approach.” (2014) Unpublished manuscript, Indiana University.
- [40] Salgado, Maria. “Choosing to Have Less Choice.” Fondazione Eni Enrico Mattei working paper, February 2006.
- [41] Samuelson, William, and Richard Zeckhauser. “Status Quo Bias in Decision Making.” *Journal of Risk and Uncertainty* 1 (1988): 7-59.
- [42] Scheibehenne, Benjamin, Rainer Greifeneder, and Peter M. Todd. “Can There Ever Be Too Many Options? A Meta Analytic Review of Choice Overload.” *Journal of Consumer Research* 37.3 (2010): 409-425.

- [43] Schram, Arthur and Joep Sonnemans. “How Individuals Choose Health Insurance: An Experimental Analysis.” *European Economic Review* 55 (2011) 799-819.
- [44] Simon, Herbert A. “Models of Man; Social and Rational.” (1957).
- [45] Sims, Christopher, “Implications of Rational Inattention.” *Journal of Monetary Economics*, 50(3)
- [46] Sonsino, Doron, Uri Benzion, and Galit Mador. “The Complexity Effects on Choice with Uncertainty – Experimental Evidence.” *The Economic Journal* 112.482 (2002): 936-965.
- [47] Taubinsky, Dmitry, and Alex Rees-Jones (2016). “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment,” NBER Working Paper No. 22545.
- [48] Tse, Alan, Lana Friesen, and Kenan Kalayci. “Complexity and Asset Legitimacy in Retirement Investment.” (2014) Working paper, University of Queensland.
- [49] Tversky, Amos, and Eldar Shafir. ”Choice under conflict: The dynamics of deferred decision.” *Psychological science* 3.6 (1992): 358-361.
- [50] Wilcox, Nathaniel T. “Lottery Choice: Incentives, Complexity and Decision Time.” *Economic Journal* 103.421 (1993): 1397-1417.

6 Appendix - Structural Estimation

We consider decision problems in which the participant’s choice y is either to invest in a portfolio ($y = 1, \dots, N - 1$) or to opt out ($y = N$). Equation (1) from above, due to Matějka and McKay (2015), implies that:

$$\begin{aligned}
 \Pr(y < N) &= \frac{\sum_{j=1}^{N-1} e^{(v_j + \alpha_j)/\lambda}}{\sum_{j=1}^N e^{(v_j + \alpha_j)/\lambda}} \\
 \Pr(y = N) &= \frac{e^{(v_N + \alpha_N)/\lambda}}{\sum_{j=1}^N e^{(v_j + \alpha_j)/\lambda}}
 \end{aligned} \tag{8}$$

We assume that participants have homothetic preferences, such that:

$$\begin{aligned}
 v_i &= m * \bar{v}_i \quad \forall i < N \\
 v_N &= \omega * \bar{v}_N,
 \end{aligned} \tag{9}$$

where \bar{v}_i is the fundamental value of option i when the endowment is \$1, \bar{v}_N is the fundamental value of opting out when the outside option is \$1, and m and ω are respectively the endowment and the outside option.

We rewrite (8) by replacing v_i with $\ln v_i$ and dividing both the numerator and denominator by $e^{(\ln \omega + \ln \bar{v}_N + \alpha_N)/\lambda}$.

$$\begin{aligned}\Pr(y < N) &= \frac{e^{(\mu + \frac{\ln m}{\lambda} - \frac{\ln \omega}{\lambda})}}{1 + e^{(\mu + \frac{\ln m}{\lambda} - \frac{\ln \omega}{\lambda})}} \\ \Pr(y = N) &= \frac{1}{1 + e^{(\mu + \frac{\ln m}{\lambda} - \frac{\ln \omega}{\lambda})}}\end{aligned}\tag{10}$$

where

$$\mu = \ln \left\{ \sum_{j=1}^{N-1} e^{(\ln \bar{v}_j - \ln \bar{v}_N + \alpha_j - \alpha_N)/\lambda} \right\}.$$

In the estimation, μ is a fixed effect in a conditional logit. μ is specific to a given slope of the budget line.

We allow the half of the sample with lower decision-making skills to have a different lambda from the other half of the sample with higher decision-making skills. Accordingly, the μ for a given slope of the budget line is allowed to vary with decision-making skill.