

Risk Aversion in a Dynamic Asset Allocation Experiment

Isabelle Brocas, Juan D. Carrillo, Aleksandar Giga, and Fernando Zapatero*

JEL Classification: C91, D03, D81, G02, G11.

Keywords: laboratory experiments, portfolio allocation, risk aversion, HARA, CRRA.

*Brocas, brocas@usc.edu, University of Southern California and Centre for Economic Policy Research (CEPR); Carrillo, juandc@usc.edu, University of Southern California and CEPR; Giga, aleksang@usc.edu, University of Southern California; Zapatero (corresponding coauthor), FZapatero@marshall.usc.edu, Marshall School of Business, University of Southern California, Los Angeles, CA 90089. The authors thank Elena Asparouhova, Glenn Harrison, Alex Imas and the members of the Los Angeles Behavioral Economics Laboratory (LABEL); participants (and especially discussants) of the 2015 Western Finance Association meetings (Thomas Gilbert) 2015; CEAR/MRIC Workshop at the Georgia State University, 2014; Economic Science Association meeting, and the seminar at USC for their helpful comments and suggestions. We thank the LUSK Center for Real Estate for their financial support.

Abstract

We conduct a controlled laboratory experiment in the spirit of Merton (1971), in which subjects dynamically choose their portfolio allocation between a risk-free and risky asset. Using the optimal allocation of an investor with hyperbolic absolute risk aversion (HARA) utility, we fit the experimental choices to characterize the risk profile of our participants. Despite substantial heterogeneity, decreasing absolute risk aversion and increasing relative risk aversion are the predominant types. We also find some evidence of increased risk taking after a gain. Finally, the session level risk attitudes show a different profile than the individual descriptions of risk attitudes.

I. Introduction

Economic models, in general, and finance models, in particular, often start with the assumption that the optimal decisions of a representative agent are a good description of the function of the economy in the aggregate. Even if the model considers multiple agents with different preferences, the preferences are assumed to belong to a narrow class and differ on the value of a single parameter. In any case, the choice of utility function of the representative agent or class of utility functions, in the case of multiple agents, is typically justified on the grounds of tractability.

A large body of work has studied individual behavior and proposed utility representations to use in models. In macroeconomics and finance, the typical methodology is to postulate a specific type of utility, derive some predictions resulting from partial or general equilibrium considerations based on that utility (or class of utilities) and, finally, empirically test the predictions.

A parallel line of analysis of this problem has been undertaken in the behavioral economics literature. This literature has mostly relied on a different methodology, namely, laboratory experiments. More precisely, this literature studies individual decision-making in narrowly defined situations. The evidence collected permits the derivation by induction of the properties a utility function should display or even piecing together of a functional form, for example, the value function of Prospect Theory (Kahneman and Tversky (1979)).

Over the last two decades, other studies related to the characteristics of preferences have used this experimental methodology; we later review the main contributions. In this paper, we propose directly estimating individual utility in an experimental setting. We elicit the functional form that best represents the decisions of participants. However, for

this approach to be practical, we must select a class of parametric utility functions that can be fitted to the data. For reasons explained below, we work with the class of *hyperbolic absolute risk aversion* (HARA) utility functions studied in Merton (1971).

The seminal work of Merton (1971) considers the class of HARA utility functions over intertemporal consumption or final wealth to study the problem of portfolio optimization of a risk averse individual investor. The HARA class is broad and nests utility functions that are used not only in the portfolio optimization literature but also in most of the asset pricing, corporate finance and macroeconomics literature. In particular, the constant relative risk aversion (CRRA) and constant absolute risk aversion (CARA) utilities are special cases of HARA.

Based on the previous considerations, we conduct an experiment that replicates Merton's (1971) setting within the technical limitations of our laboratory environment. Subjects choose how to invest their wealth between two assets, one safe and one risky, over the course of 15 investment paths. The subject (she) starts a path with an initial endowment, which she allocates between the assets. After observing the returns of the assets, she reallocates her wealth, and new returns are observed. This dynamic process lasts for 10 periods, after which her final payoff is recorded, and a new path is started with the same initial endowment as the previous path. Overall, each subject makes 150 investment decisions with different levels of wealth.

We must emphasize that the HARA class does not include all types of utilities discussed in the literature. In particular, the HARA utility assumes risk-aversion, while the value function of Prospect Theory allows risk-loving, which can explain decisions observed both in experimental settings and in practical situations inconsistent with HARA utilities. However, we observe a low level of risk-loving behavior in our data (in less than 11% of the

subjects). Other types of preferences, such as habit formation (Campbell and Cochrane (1999)) or recursive preferences (Epstein and Zin (1989)), are also not included in the HARA class. On the other hand, Merton (1971) provides a tractable dynamic setting that is suitable for an experimental setting, which allows us to assess whether many of the utilities used in economic models, namely, CRRA and CARA, are consistent with individual decisions as well as whether other utilities within the HARA class provide a better representation. Another important observation is that HARA utilities do not aggregate in general. That is, even if we could corroborate that all economic agents display HARA preferences, we would not be able to conclude that it is possible to construct a representative agent, with the exception if all agents belong to some narrow subset (e.g. CRRA). Our experiment provides a tool that characterizes a broader set of risk attitudes, which in turn can be used as a guide to determine the preferences (if any) of the representative agent.

This experimental setting allows us to address the following specific questions. First, consistent with the experimental literature on preferences, do we observe substantial heterogeneity in the risk attitudes of our subjects? Second, can we fit the data well using a structural estimation of an expected HARA utility model? If so, what type of HARA utility best explains the investment strategy of each participant? Third and related, do we observe frequent and/or severe deviations from neoclassical theory, i.e., systematic biases at odds with standard expected utility theory? Fourth, if we analyze the data at the session level, how does the group behavior compare to that of subjects considered individually?

Our starting point is the optimal portfolio allocation of an expected HARA utility maximizer, as derived in Merton (1971). Given this analytical characterization, we estimate the absolute and relative risk aversion parameters of our subjects using the 150 choices

made in the experiment.

Our main findings are as follows. Consistent with the existing literature, our experimental subjects (undergraduate students) are highly heterogeneous. At the same time, some risk attitudes are more prevalent than others. Most individuals increase the total amount of wealth invested in the risky asset as their wealth increases (decreasing absolute risk aversion or DARA). They also decrease the fraction of wealth invested in the risky asset as their wealth increases (increasing relative risk aversion or IRRA). Overall, more than half of our subjects can be confidently classified in the combined DARA-IRRA category, the risk attitude conjectured by Arrow (1971) to be the most natural among investors.

We also find some evidence of biases that is inconsistent with the assumptions of standard expected utility theory. Some subjects (19%) change their risk-taking behavior over time. More significantly, 44% of subjects exhibit a gain/loss asymmetry. Of these, the vast majority (39%) take more risks after a gain, while only 5% take more risks after a loss. Overall, many subjects exhibit some type of anomaly relative to the standard expected utility theory. However, these are small in magnitude, which is why the expected utility model performs well despite their presence.

Finally, we conduct a session level analysis using two different methodologies. We find that some types (notably, CARA) are not present at the session level, even though there are such individuals among our subjects. Also, the relative risk aversion coefficient estimated for the sessions is typically lower than those of individuals. Therefore, while the risk attitudes of most individuals are best captured by DARA-IRRA, many of the aggregate parameters are consistent with the DARA-DRRA or DARA-CRRA types. This result occurs because DRRA agents accumulate, on average, more wealth than IRRA agents and therefore end up having a greater impact in the session.

Before proceeding to the analysis, we present a brief literature review. Methods to elicit risk attitudes in static settings abound in economics. Perhaps the best-known and most widely employed technique is the “list method” proposed by Holt and Laury (2002), hereafter [HL]. This method is fast, intuitive and easy to implement. The list method offers an excellent and simple measure to compare risk attitudes across individuals and has been extended in several directions either to improve the precision of estimates (Andersen, Harrison, Lau, and Rutström (2006), Maier and Ruger (2012)) or to obtain a more efficient algorithm (Wang, Filiba, and Camerer (2010)).¹ However, simplicity comes at the expense of a design that is not intended (and therefore not suitable) to provide a *precise* measure of the risk preference of individuals endowed with a general utility function. For example, the [HL] procedure assumes CRRA utility; therefore, by construction, it cannot assess the changes in the percentage of risk taking as a function of wealth. The [HL] procedure also provides only interval estimates of the parameter, so it is difficult to assess the fit of the data according to the utility specification and to challenge the model.

Risk attitudes have been explored in dynamic settings as well, most notably in the game show “Deal or No Deal”. Assuming a CRRA functional form, Post, Van den Assem, Baltussen, and Thaler (2008) find that the Expected Utility Theory cannot explain the contestants’ decisions well and point out that previous outcomes play a significant role in the choices of participants. Andersen, Harrison, Lau, and Rutström (2008) perform a laboratory replication of “Deal or No Deal”. They estimate average risk preferences

¹Other, almost equally simple, risk elicitation designs have been proposed by Becker, DeGroot, and Marschak (1964), Binswanger (1980), Hey and Orme (1994), Gneezy and Potters (1997), Eckel and Grossman (2008), and Sokol-Hessner, Hsu, Curley, Delgado, Camerer, and Phelps (2009), among others. For surveys of empirical and experimental elicitation procedures and results, we refer to Harrison and Rutström (2008), Charness, Gneezy, and Imas (2013) and Friedman, Isaac, James, and Sunder (2014).

without constraining the utility model to a single parameter and find moderate levels of risk aversion, with evidence suggesting IRRA. In Rapoport (1984) and Rapoport, Zwick, and Funk (1988), subjects invest in risky securities and a safe asset in a dynamic setting, and evidence is found in favor of IRRA and against CARA or CRRA. Recently, Levy and Levy (2017) showed that CRRA may be a good approximation for decisions facing large (albeit hypothetical) stakes. However, their experimental design does not allow for periodic portfolio revision.

Our methodology has a number of advantages over previous experimental designs. First, we can structurally estimate an asset allocation model based on a rich class of utility functions for each *individual* subject. We can also determine the loss in predictive and explanatory power when we restrict our analysis to simpler utility functions. Second, we can measure standard errors of individual estimates and assess the fit of the data. We can also study the structure of the noise and its relationship with wealth levels. Third, our dynamic framework is useful for measuring behavioral anomalies due to repeated exposure to risk. We can detect any gain/loss asymmetry in behavior and determine whether a subject changes her risk attitude over the course of the experiment.

Given the investment nature of our task, our paper also relates to market experiments in which most of such tasks are implemented.² Levy (1994) proposes a non-structural analysis to study risk attitudes in a market experiment. As in our paper, his results overwhelmingly support DARA but, unlike us, he does not find evidence in support

²Market experiments have been extensively used in finance research to analyze asset bubbles (Smith, Suchanek, and Williams (1988), Haruvy and Noussair (2006)), to test the predictions of asset pricing models (Plott and Sunder (1988); Bossaerts and Plott (2004); Bossaerts, Plott, and Zame (2007)) and to test investor behavior (Bossaerts, Ghirardato, Guarnaschelli, and Zame (2010); Frydman, Barberis, Camerer, Bossaerts, and Rangel (2014)) among other subjects.

of IRRA. Contrary to this literature, our main goal is to isolate risk attitudes, which is why we opt for an individual decision-making rather than a market set up.³ We also provide complete information about the design to prevent subjects from forming beliefs we could not observe.⁴

Finally, our results on path dependence of choices and gain/loss asymmetry are related to the literature that highlights behavioral anomalies in choice under uncertainty.⁵ Our design is not intended to test for specific behavioral anomalies nor to fit behavioral models. However, consistent with Thaler and Johnson (1990), we find that prior gains (losses) decrease (increase) risk aversion for many of our subjects.⁶

This paper is organized as follows. In section II, we present the theoretical framework. In section III, we describe the experimental setting. In section IV, we present the econometric model and results of the classification analysis and estimation. In section V, we investigate behavioral anomalies. In section VI, we provide an aggregate analysis of

³Other related individual asset allocation experiments test whether subjects allocate portfolios efficiently (Kroll, Levy, and Rapoport (1988); Kroll and Levy (1992); Sundali and Guerrero (2009)).

⁴Indeed, we spend substantial effort during the instruction period to explain, in detail, the financial environment of the experiment so that expectations play as small of a role as possible (for a survey on the rapidly expanding experimental literature studying the effect of expectations on risk taking behavior in macroeconomics and finance, we refer to Assenza, Bao, Hommes, and Massaro (2014)).

⁵Discrepancies between observed behavior and theoretical predictions may come from errors in choices (Jacobson and Petrie (2009)), reference dependent preferences (Koszegi and Rabin (2006); Abeler, Falk, Goette, and Huffman. (2011); Knetsch and Wong (2009); Ericson and Fuster (2011); Sokol-Hessner et al. (2009)), or disappointment aversion (Choi, Fisman, Gale, and Kariv (2007); Gill and Prowse (2012)) among other reasons.

⁶The way prior outcomes affect subsequent risk taking is not a settled matter. See Imas (2016) for a summary and set of experiments showing how different types of losses (paper vs. realized) produce different risk choices immediately after.

the data. In section VII, we offer some concluding remarks. An analysis of the explanatory and predictive power of our expected utility model is relegated to Appendices B and C.

II. Theory

Consider a continuous-time setting with a risk-free security that pays a constant interest rate and a single risky security whose price satisfies a geometric Brownian motion process. Merton (1971) shows that, for the class of Hyperbolic Absolute Risk Aversion (HARA) utility, the optimal investment policy of the economic agent has an explicit solution.⁷

At each instant t , an agent (she) allocates her wealth $X(t)$ between two assets, a risky asset A and safe asset B . At $t = 0$, her initial wealth is $X(0) = x_0 > 0$. The temporal horizon is finite and equal to T . The agent can reallocate her portfolio at each instant t until date T , which is the time at which she enjoys her accumulated wealth $X(T)$. Therefore, at each t , she maximizes the expected utility of wealth at time T . We assume that the agent's preferences are characterized by the general Hyperbolic Absolute Risk Aversion (HARA) utility function with the two parameters, γ and η , first used by Merton (1971) in a dynamic portfolio allocation.

Formally:

$$(1) \quad U(X) = \frac{1-\gamma}{\gamma} \left(\frac{X}{1-\gamma} + \eta \right)^\gamma$$

⁷The previous setting also amounts to dynamic completeness. This notion is studied in an experimental setting by Bossaerts, Meloso, and Zame (2008).

with the following restrictions:

$$\gamma \neq 1, \quad \frac{X}{1-\gamma} + \eta > 0 \quad \text{and} \quad \eta = 1 \text{ if } \gamma = -\infty$$

This family of utility functions is rich in the sense that it encompasses utility functions with absolute and relative risk aversion that are increasing, constant or decreasing depending on γ and η .⁸ The agent exhibits decreasing absolute risk aversion when $-\infty < \gamma < 1$ and constant absolute risk aversion when $\gamma \rightarrow +\infty$ or $\gamma \rightarrow -\infty$. She exhibits increasing, constant and decreasing relative risk aversion when $\eta > 0$, $\eta = 0$ and $\eta < 0$, respectively.

The price of the safe asset $B(t)$ evolves as follows:

$$(2) \quad dB(t) = rB(t)dt$$

where $r > 0$. The price of the risky asset $A(t)$ follows a geometric Brownian motion process with drift μ ($> r$) and diffusion σ (> 0). Formally:

$$(3) \quad dA(t) = \mu A(t)dt + \sigma A(t)dW(t)$$

where $W(t)$ is a standard Brownian motion process. Let $\pi(t)$ be the amount of wealth

⁸A more general specification of the HARA utility function is: $U(X) = \frac{1-\gamma}{\gamma} \left(\frac{\beta X}{1-\gamma} + \eta \right)^\gamma$. In our case, the parameter β is not identified and cannot be estimated.

allocated to the risky asset A at date t . The wealth $X(t)$ grows as follows:

$$\begin{aligned} dX(t) &= \pi(t)\mu dt + \pi(t)\sigma dW(t) + [X(t) - \pi(t)]r dt \\ &= [X(t)r + \pi(t)(\mu - r)] dt + \pi(t)\sigma dW(t) \end{aligned}$$

At each date t , the agent solves the following problem \mathcal{P} :

$$\begin{aligned} \mathcal{P} : \quad & \max_{\pi} E[U(X(T))] \\ \text{s.t.} \quad & dX(t) = [X(t)r + \pi(\mu - r)]dt + \pi(t)\sigma dW(t) \\ & X(0) = x_o \end{aligned}$$

Given the complete markets assumption and specification of utility and asset returns, our problem has a closed-form solution that we summarize in the next result.

Proposition 1 *If markets are complete and time is continuous, the optimal amount allocated to the risky asset at date t when the accumulated wealth is $X(t)$ is:*

$$(4) \quad \hat{\pi}(t) = \frac{\mu - r}{\sigma^2} \left(\frac{X(t)}{1 - \gamma} + \eta e^{-r(T-t)} \right)$$

Proof. It is straightforward from the by now standard martingale representation methodology of Karatzas, Lehoczky, and Shreve (1987) and Cox and Huang (1989).

The amount allocated to the risky asset depends on the current wealth $X(t)$, the investment horizon left $T - t$, the parameters that characterize the return dynamics of the assets, and the risk aversion parameters. The model predicts that the amount allocated to the risky asset increases in the current wealth if the agent exhibits decreasing absolute risk aversion ($\gamma < 1$). Also, the allocation depends on current wealth irrespective of how wealth

has been accumulated in the past. Finally, when $\eta > 0$ (respectively, $\eta < 0$), $\hat{\pi}(t)$ increases (respectively, decreases) as time passes. Note that $\eta = 0$ corresponds to the CRRA specification, where the agent invests a constant proportion of her wealth in the risky asset irrespective of the level of wealth and the horizon left to invest.

The risk attitude of each agent is characterized by two dimensions, absolute risk aversion (ARA) and relative risk aversion (RRA), which can each be increasing (I), constant (C) or decreasing (D) in wealth. These dimensions are determined by the (γ, η) parameter combination of the individual, which we call "type". Equation (4) predicts each type in terms of the amount of wealth invested in the risky asset. First, all types with DARA increase the risky investment as wealth increases. Of these, an agent with decreasing relative risk aversion (DARA-DRRA type) is willing to short-sell when her wealth is low ($\hat{\pi}(t) < 0$ when $X(t)$ is small). By contrast, an agent with increasing relative risk aversion (DARA-IRRA type) is willing to borrow when her wealth is low ($\hat{\pi}(t) > X(t)$ when $X(t)$ is small). Second, types with IARA decrease the risky investment as wealth increases. Of these, an agent with increasing relative risk aversion (IARA-IRRA type) will invest a positive amount of wealth in the risky asset only when her wealth is low.

The closed-form solution for the optimal portfolio requires complete markets. Complete markets allow investors to borrow and take short positions in the risky security which, given the nature of the experiment, we must rule out. Therefore, our results are an approximation. When we present our results, we will discuss the impact of these restrictions. Nevertheless, some of the qualitative properties of the solution are not affected by them. In particular, agents represented by DARA, CARA and IARA utility functions will, respectively, choose to invest (weakly) more, the same and (weakly) less total amounts in the risky asset as their wealth increases. Similarly, agents represented by DRRA, CRRA

and IRRA utility functions will, respectively, choose to invest a (weakly) larger, equal and (weakly) smaller fraction of their wealth in the risky asset as their wealth increases.

Table 1. Risk Types as a Function of Risk Aversion Parameters.

	$\gamma < 1$	$\gamma > 1$	$\gamma = -\infty$	$\gamma = +\infty$
$\eta < 0$	DARA-DRRA*	—	—	—
$\eta = 0$	DARA-CRRA	—	—	—
$\eta > 0$	DARA-IRRA**	IARA-IRRA**	CARA-IRRA**	CARA-IRRA**
	$\partial\hat{\pi}/\partial X > 0$	$\partial\hat{\pi}/\partial X < 0$	$\partial\hat{\pi}/\partial X = 0$	$\partial\hat{\pi}/\partial X = 0$
* $\hat{\pi} = 0$ for small X and $\hat{\pi} = X$ for large X ; ** $\hat{\pi} = X$ for small X and $\hat{\pi} = 0$ for large X				

Table 1 summarizes the risk types as a function of η and γ , given the parametric restrictions in the utility function. It shows which types are likely to be constrained and for which wealth levels (indicated by * and **) and also shows how investment varies with wealth.

III. Experimental Design

The main objective of this paper is to study the dynamic portfolio choice of agents in a controlled laboratory setting. To this purpose, we design a dynamic investment problem that follows as closely as technically feasible the setting of the theory section. Subjects in the experiment allocate wealth between one safe and one risky asset during 15 investments *paths* consisting of 10 *periods* each. The experiment consists of 13 sessions run in the Los Angeles Behavioral Economics Laboratory (LABEL) at the University of Southern California.⁹ Each session has between 7 and 10 subjects for a total of 120 recruited subjects, of which 3 are omitted from the analysis due to software malfunction. All subjects participate in three treatments that are always performed in the same order.

⁹For information about the laboratory, please visit <http://dornsife.usc.edu/label>.

The first treatment corresponds to the paradigm under study in this paper. The results of the other two treatments are reported in Brocas, Carrillo, Giga, and Zapatero (2016).

Each subject starts each path in period 1 with an endowment of \$3, which she allocates between two assets: a risky asset A and safe asset B . After period 1 ends, each subject earns a return on her portfolio and moves to period 2. She then reallocates her portfolio and earns new returns. This process continues for a total of 10 periods. After period 10, the investment path ends and the subject's final payoff in that path is recorded. Each subject then moves to the next investment path, where her endowment is reset to \$3. Subjects have 10 seconds to make their decision in period 1 of each path and 6 seconds in periods 2 to 10. They all begin and end their investment paths at the same time. All subjects go through 15 paths for a total of 150 choices. Subjects know at the beginning of the experiment the number of paths and periods in each path they will go through.

The return of the safe asset B is 3%, while the return of asset A is drawn from a log-normal distribution with a mean of 23.5% and standard deviation of 73.4%.¹⁰ The

¹⁰This (unrealistically high) mean and standard deviation ensure enough volatility in returns for generating interesting wealth effects and comparative statics. In the discrete version of the experiment, the evolution of wealth is:

$$X(t+1) = X_B(t)(1+r) + X_A(t)e^R$$

where $X_i(t)$ is the dollar amount invested in asset $i \in \{A, B\}$ and R normally distributed with a mean of 0.06 and standard deviation of 0.55. Note that the return of the risky asset, e^R , is log-normally distributed, so the worst case for the subject is to lose her investment in A . In a part of the instructions and on the upper left corner of the screen, we described in words the parameters of the return on asset A as being normally distributed with a mean of 6% and standard deviation of 55%, when it should have read log normally distributed with a mean of 23.5% and a standard deviation of 73.4%. In other words, we accidentally described R instead of e^R . Nevertheless, we are confident that this did not impact the results, as the rest of the instructions and accompanying slides vividly and correctly describe the entire distribution of Asset A through graphical and video examples. Moreover, students are shown a correctly specified interactive

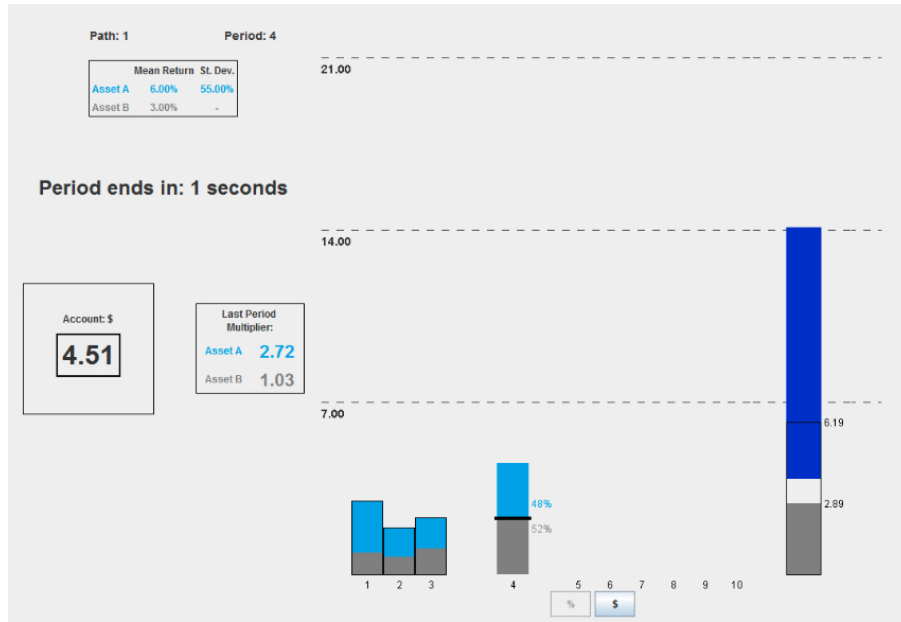
parameters do not change throughout the experiment. The draw of the return is presented in the form of a multiplier, that is, the number that multiplies the allocation to that asset. Importantly, all participants in a session are subject to the same draws, which makes it possible to analyze the aggregate portfolio of each session (see section VI). At the same time, we make clear to each subject that her return is in no way affected by the allocation decision of the other subject.

Figure 1 provides a screenshot that describes what a subject sees in a given period of a path. Current wealth is represented by the vertical bar above the current period number (period 4 in this example). Initially, the bar is not active and wealth is not allocated to either asset. Subjects must click on the bar to activate wealth and move a horizontal slider to divide their wealth between assets A and B . The upper portion of the bar represents the money invested in A and the lower portion represents the money invested in B . The figures on the right side of the bar show the allocation. To facilitate her reasoning, each subject may change the display of the allocation at any time between the percentage invested in each asset (box labeled “ % ”) and the total amount in each asset (box labeled “ \$ ”). After the period expires, returns are applied and subjects move to the next period. A new bar with a height corresponding to the new wealth appears to the right of the previous one for the new period and becomes inactive again. Subjects must reactivate it to choose a new allocation; otherwise, they earn no interest in that period and their account simply carries

projection bar at the end of the screen that informs them of the possible distribution of payoffs at the end of the path as they change their current allocation. Previous research shows the importance of visual and interactive tools for financial literacy (Lusardi, Samek, Kapteyn, Glinert, Hung, and Heinberg (2017)). Lastly, students had five practice paths and a quiz before starting their paid trials, which was enough to experiment with the bar and the payoffs.

over.¹¹ Subjects observe bars to the left of the current one that remind them of their past allocations and returns. These bars accumulate up to period 10 and then are reset for the new path. Finally, the left-hand side of the screen shows a summary of the information of the main ingredients of the experiment: (i) the current path and period; (ii) a reminder of the mean and standard deviation of returns of assets A and B ; (iii) the time remaining to make a choice in the current period; (iv) the accumulated wealth in the current path; and (v) the multiplier of assets A and B in the last period of the current path.

Figure 1. Screenshot of Path 1 - Period 4.



This dynamic wealth allocation problem is challenging and may require substantial learning. We develop a highly illustrative 40-minute instruction period using numerical examples, videos, five practice paths and a comprehension quiz (instructions can be found in Appendix A). To help with the cognitive strain, we also add a projection bar to the right side of the screen that tells the subject what she would expect if she were to keep her

¹¹This helps prevent subjects' inertia and a bias towards any status quo allocation. Level of inactivity in our experiment was negligible.

current investment strategy until the last period. The bar shows the potential accumulated earnings from asset B and identifies the 20th, 50th and 80th percentile of the earning distribution from asset A (see Figure 1). As the participant changes her allocation the projection bar automatically adjusts.¹²

Each participant received a \$5 show-up fee and her final earnings in the final period of *one* randomly selected path (the average earnings were \$9.5 with a maximum of \$41)¹³. At the end of the experiment, we collected answers to education, demographic and income related questions as well as their own description of the strategies employed. The length of the experiment, including all three treatments and the survey, was two hours.

Note that the experimental design closely follows the theory with two important differences, both of which were introduced for technical reasons. First, choices are made in discrete time, with only 10 decisions per path. Continuous time is difficult to implement in an experimental setting (although not impossible, see e.g. Friedman and Oprea (2012))¹⁴. Second, we do not allow our participants to borrow or short sell, which means that the markets are incomplete in the sense of Merton (1971). Borrowing and short selling are difficult to implement experimentally since they may result in taking money away from participants. Our data analysis takes this restriction into account.

¹²We carefully explain the function of the bar by simulating a large number of period-by-period trajectories of wealth coming from a given allocation strategy.

¹³Subjects were also compensated for the other two treatments. The total compensation in the experiment averaged \$23, with a maximum of \$244.

¹⁴Duffie and Protter (1992) provide a theoretical discussion on the convergence of discrete-time processes to continuous-time ones.

IV. Results

Our first objective is to test how well the expected utility theory fits the data. We adopt a structural approach and estimate the risk parameters (γ, η) of each subject assuming they behave according to the expected utility theory model. This approach is used to classify our subjects according to their risk type.

A. Econometric Model

According to equation (4) and subject to the above-mentioned caveats of incomplete markets and discrete time, the expected utility theory predicts that the portfolio allocation and wealth will vary over time according to the following system:

$$\begin{cases} \hat{\pi}(t) = \frac{\mu - r}{\sigma^2} \left(\frac{X(t)}{1 - \gamma} + \eta e^{-r(T-t)} \right) \\ dX(t) = [X(t)r + \hat{\pi}(t)(\mu - r)] dt + \hat{\pi}(t)\sigma dW(t) \end{cases}$$

The parameters γ and η can be estimated from the first equation using least squares fitting. Since our data are obtained in discrete time, we consider the discrete version of the model. For each individual, in each path i and at each period t we observe the current wealth $X_{i,t}$ and the chosen allocation of this wealth to the risky asset $\pi_{i,t}$. Let $F_t = e^{-r(T-t)}$, our structural econometric model given HARA utility is $\mathcal{M}^{\text{HARA}}$:

$$(5) \quad \pi_{i,t} = aX_{i,t} + bF_t + u_{i,t}$$

where $a = \frac{\mu-r}{\sigma^2(1-\gamma)}$, $b = \frac{(\mu-r)\eta}{\sigma^2}$ and $u_{i,t} \sim \mathcal{N}(0, \sigma_u^2)$ is an error term.¹⁵ Given a and b , the parameters γ and η are identified. In the next section, we classify the risk attitude of our subjects by fitting this model to their decisions.

Note that a myopic decision-maker would maximize the instantaneous expected utility $E[U(X(t))]$ at each period t . This problem has a simple closed-form solution: the optimal allocation in the risky asset is obtained by replacing $e^{-r(T-t)}$ with 1 in the equilibrium equation of Proposition 1. For our data, $e^{-r(T-t)} \in [0.7, 1]$. This value is close enough to 1 to make the myopic model very similar to the forward-looking model.¹⁶

Also, we require enough variation in wealth within subjects for an accurate estimation. In half of the sample, the 5th and 95th percentiles of wealth are approximately \$1 and \$15, respectively. For the other half of the sample, the range extends from \$1 to \$20, respectively. Although these figures are not excessively large, the dispersion is important enough to obtain reliable estimates of absolute and relative risk aversion.

Lastly, our structural model is well specified only if subjects do not systematically invest all their wealth in the safe or the risky asset, which poses a challenge. On one hand, treating the data as if all choices are interior biases the interpretation of the parameters and the residuals of the regression. On the other hand, eliminating the constrained choices from the analysis also biases the estimated parameters. The solution we propose is to separately classify subjects who hit the bounds often from those who do not.

¹⁵We relax the assumptions on the error term's distribution later (see subsection C.1). For robustness purposes, we evaluate the data from unconstrained subjects using the beta regression model. We also test the sensitivity of our risk elicitation to the subject's perception of the return parameters. In both analyses, available from the authors upon request, the results remain qualitatively the same.

¹⁶We conducted the analysis based on the myopic model and did not find any qualitative changes in the classification of our subjects.

B. Classification Criteria: Constrained vs. Unconstrained

Subjects

Our first task is to empirically determine which subjects are affected by the inability to short-sell (i.e, to set $\pi_t < 0$) and/or borrow (i.e, to set $\pi_t > X_t$). For the large majority of our subjects, the pressure to short-sell or borrow is low. At the aggregate level, subjects invest all their wealth in the safe asset 2.2% of the time and in the risky asset 8.2% of the time.¹⁷ At the individual level, there is heterogeneity in behavior. Table 2 shows the distribution of subjects as a function of their likelihood to hit the constraints.

Table 2. Number of Subjects Affected by the Constraints.

	% trials			Total
	(0%, 10%]	(10%, 20%]	(20%, 100%]	
Hit $\pi_t = 0$ only	10	0	0	10
Hit $\pi_t = X_t$ only	24	3	8	35
Hit $\pi_t \in \{0, X_t\}$	13	11	14	38
$\pi_t \in (0, X_t)$ always	n/a	n/a	n/a	34

Only 34 subjects never hit a constraint. However, if we combine these subjects with those who hit the constraints no more than 10% of the time, we account for 81 individuals, or 69% of the sample. We call these subjects ‘unconstrained’. Of the remaining subjects, 11 would have liked to borrow and 25 would have liked to both borrow and short sell. We call these subjects ‘constrained’.

¹⁷A choice is defined as non-constrained (interior) when the allocation to the risky asset is bigger than 2% and smaller than 98% of the wealth.

C. Estimation and Classification

C.1. Unconstrained Subjects

We estimate the risk aversion parameters (γ, η) of the 81 unconstrained subjects for which the econometric model $\mathcal{M}^{\text{HARA}}$ is well specified. Given that our observations are repeated measures for the same subject and that wealth follows a stochastic process, we must be careful about issues that arise naturally in this time series framework and that may contradict the underlying assumptions required to use the least squares method.

First, the error term should have a constant variance. We run a standard OLS on each individual's dataset and apply the White test to detect the presence of heteroscedasticity. We find that the variance of the residuals increases with the level of wealth for 73 out of the 81 unconstrained subjects (at the 5% significance level) and is constant for the rest.

Second, error terms should be uncorrelated across periods. We test for serial correlation for each participant by looking at the residuals of the OLS regression, denoted by $\hat{u}_{i,t}$. Note first that an error at period $t - 1$ applied to the amount invested in the risky asset at that period affects the wealth level at period t . Therefore, regressors are not independent of the error term. To account for this, we use the Breusch-Godfrey test, which allows explanatory variables to not be strictly exogenous. Formally, we consider the regression:

$$\hat{u}_{i,t} = \beta_0 + \beta_1 X_{i,t} + \beta_2 F_t + \rho \hat{u}_{i,t-1} + v_{i,t}$$

where $X_{i,t}$ and F_t account for weak exogeneity and $v_{i,t}$ are assumed to be i.i.d. with normal distribution $\mathcal{N}(0, \sigma_v^2)$. We use robust standard errors in our test and find first-order serial correlation ($\rho > 0$) for 63 out of 81 subjects. To correct for heteroscedasticity and

autocorrelation, we run the OLS regression with Newey-West standard errors.

Figure 2 displays the estimated (γ, η) risk parameters of the 81 unconstrained subjects using the structural model $\mathcal{M}^{\text{HARA}}$ presented in equation (5). Table 3 reports the relative and absolute risk aversion attitudes based on the estimated parameters.

Figure 2. Estimated Parameters of the 81 Unconstrained Subjects

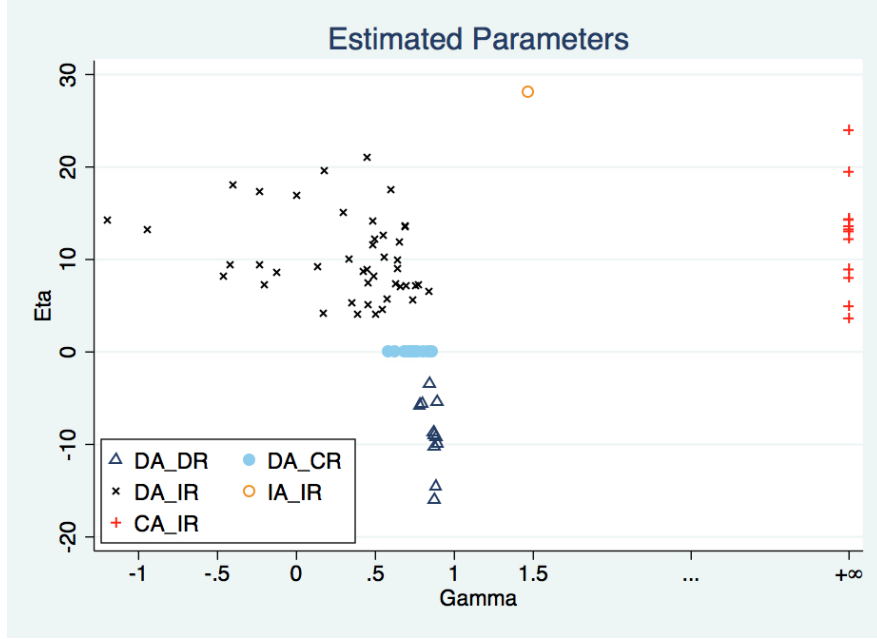


Table 3. Risk Attitude of the Unconstrained Subjects.

Risk attitude	No. of subjects
DARA-DRRA	11
DARA-CRRA	13
DARA-IRRA	44
IARA-IRRA	1
CARA-IRRA	12
Total	81

We observe substantial heterogeneity in risk attitudes. At the same time, the vast majority of subjects are DARA ($\gamma < 1$ for 84% of subjects) and IRRA ($\eta > 0$ for 70% of subjects). Overall, 54% of subjects are willing to increase their total investment in the risky asset and decrease the fraction of investment in the risky asset as their wealth increases. These are the DARA-IRRA subjects ($\gamma < 1$ and $\eta > 0$) conjectured by Arrow (1971) to

empirically be the most plausible types. By contrast, the simple one parameter specifications commonly used in the literature do not capture the risk attitude of many of our subjects well: only 15% of our subjects are CARA ($\gamma \rightarrow +\infty$) and 16% are CRRA ($\eta = 0$).¹⁸

We next compare our results to existing estimates in the literature, such as, [HL]. For this, we estimate our structural model assuming the familiar functional form:

$$\tilde{U}(X) = \frac{X^{1-\xi}}{1-\xi}$$

used in [HL]. In this case, the solution to the problem \mathcal{P} described in section II is well-known. Indeed, the agent invests a constant fraction of wealth in the risky asset:

$$\hat{\pi}(t) = \frac{1}{\xi} \frac{\mu - r}{\sigma^2} X(t).$$

Analogously to our strategy in section A, we estimate ξ from the following econometric model:

$$(6) \quad \pi_{i,t} = cX_{i,t} + \nu_{i,t}$$

where $c = \frac{\mu-r}{\xi\sigma^2}$ and $\nu_{i,t} \sim \mathcal{N}(0, \sigma_\nu^2)$ is an error term. We call this model $\mathcal{M}^{\text{CRRA}}$ and compare our estimates of ξ to those in [HL]. Because of the way the experiment is designed, [HL] only gives range estimates for the parameter ξ . Table 4 reports the proportion of

¹⁸By $\eta = 0$ we mean that the estimated parameter is not statistically different from zero at the 95% confidence interval. For our classification, we use CARA and CRRA as the null hypotheses which may over-classify subjects in those categories.

subjects who fall in each range of ξ in our model ($\mathcal{M}^{\text{CRRA}}$) as well as in the low stakes (\$0.10 to \$3.85, HL-low) and high stakes (\$2 to \$77, HL-high) treatments of [HL].

Table 4. Risk Estimates: a Comparative Analysis with Holt and Laury (2002)

Risk Aversion	HL-low	HL-high	$\mathcal{M}^{\text{CRRA}}$
$\xi < 0.15$.34	.19	.11
$0.15 \leq \xi < 0.41$.26	.19	.61
$0.41 \leq \xi < 0.68$.23	.23	.27
$0.68 \leq \xi < 0.97$.13	.22	.01
$0.97 \leq \xi < 1.37$.03	.11	.00
$1.37 \leq \xi$.01	.06	.00
no. of subjects	175	150	81

Our estimates are substantially more concentrated than in [HL]. Only 11% of our subjects exhibit risk-neutrality or risk-loving preferences ($\xi < 0.15$) as opposed to 34% and 19% in [HL-low] and [HL-high], respectively. Unlike [HL], we also find no evidence of high ($0.97 \leq \xi < 1.37$) or extremely high ($1.37 \leq \xi$) risk aversion. Overall, we have twice as many subjects as [HL] in the expected range (88% against 49% and 42% in $0.15 \leq \xi < 0.68$). These differences are important and are partly due to differences in the design and partly due to the misspecification of the CRRA utility function in our experiment (and possibly in theirs as well). These differences highlight the advantages of a rich experimental setting to better estimate risk aversion and a two-parameter specification to capture the heterogeneity present in the relative risk aversion of subjects.

C.2. Constrained Subjects

Next, we study the risk attitude of the 36 subjects who, according to the analysis in section B, are constrained by their inability to borrow and short sell. As noted before, the tendency to invest all wealth in the safe or the risky asset should depend on the amount of

wealth. We first assess how wealth affects their probability of hitting each bound. More specifically, we estimate a probit regression on the following two models:

$$\pi_{i,t}^{max} = b_0^{max} + b_1^{max}w_{i,t} + \epsilon_{i,t}^{max}$$

$$\pi_{i,t}^{min} = b_0^{min} + b_1^{min}w_{i,t} + \epsilon_{i,t}^{min}$$

where $\pi_{i,t}^{max}$ takes a value of 1 if $\pi_{i,t} = w_{i,t}$ and 0 otherwise and $\pi_{i,t}^{min}$ takes a value of 1 if $\pi_{i,t} = 0$ and 0 otherwise. We establish an effect when b_1^{max} or b_1^{min} are different from zero at the 5% significance level.

We find three distinct groups of individuals. There are 28 “constrained IRRA” subjects, who invest their entire wealth in the risky asset when their wealth is low enough ($b_1^{max} < 0$), their entire wealth in the safe asset when their wealth is high enough ($b_1^{min} > 0$), or both. This behavior is consistent with IRRA, although it can also be compatible with risk neutrality for low enough wealth levels.¹⁹ There is 1 “constrained DARA-DRRA” subject who invests his entire wealth in the safe asset when his wealth is low enough ($b_1^{min} < 0$) and in the risky asset when his wealth is high enough ($b_1^{max} > 0$). This behavior is consistent only with DARA-DRRA. Finally, there are 7 “constrained irregular” subjects who exhibit an irregular and volatile behavior with no discernible patterns or statistically significant effects. The result (which is the analogue of Table 3 for the constrained subject sample) is summarized in Table 5.

Overall, as for the unconstrained subjects, there is substantial heterogeneity among the constrained subjects. A majority (78%) exhibit increasing relative risk aversion, and

¹⁹Of these subjects, 15 are best classified as DARA-IRRA, 10 are best classified as CARA-IRRA, and 3 are best classified as IARA-IRRA.

Table 5. Risk Attitude of Constrained Subjects.

Risk attitude	No. of subjects
Constrained IRRA	28
Constrained DARA-DRRA	1
Constrained irregulars	7
Total	36

almost half (44%) exhibit decreasing absolute risk aversion.

In Appendix B, we present a number of robustness checks. We show that the class of HARA utility functions explains the investment decisions of the participants well and has good predictive power in out-of-sample analysis. We find a statistically reliable relationship between the investment decision and the set of independent variables. Also, when we estimate the parameters using a subsample (either eight randomly chosen paths or the first five periods of all paths), we can predict the behavior in the complementary subsample well. Finally, if we restrict attention to CRRA utility, the accuracy of our overall estimates suffers in a statistically significant way for one-third of our subjects. We conclude that CRRA utility has appeal due to its simplicity and analytical properties. However, it might come at the cost of a bias for a sizeable proportion of the population.

Finally, we use our questionnaire to study the correlation between risk attitude and demographics. We find an over-representation of males in the population of subjects who are affected by the inability to borrow (hit $\pi_t = X_t$). More precisely, 33% of males vs. 18% of females are in the Constrained IRRA group. Among the unconstrained subjects, the distributions of types in the male and female populations are not significantly different.

V. Behavioral Anomalies

Several studies have reported behavioral anomalies in decision-making under risk and uncertainty. One notable anomaly is the tendency of subjects to repeat choices that have generated gains in the past and avoid choices that have generated losses in the past. In a financial setting, this tendency translates into repeating risky investments after a gain and moving wealth into safe assets after a loss, even when draws are known to be i.i.d. (Thaler and Johnson (1990)).

A second and related anomaly is a disproportionate preference to avoid losses relative to acquire gains, as in prospect theory (Kahneman and Tversky (1979)). In a financial setting, the reference point can be the current wealth or any other heuristic. From a dynamic perspective, the reference point is likely to change over time, suggesting that a certain degree of time dependence may be observed.²⁰

Our goal in this section is to determine if there are systematic biases in choices due to dynamic considerations rather than to test specific models or fit specific parametric functions. In our dynamic expected utility model, negative or positive shocks at $t - 1$ affect wealth at t and therefore the investment decision at t . The risk attitude of each subject determines how she should respond to positive or negative shocks. To test whether subjects react differently after a positive and negative shock or whether time dependence is present, we must control for any effect that emerges naturally from the model. To do so, we study the residuals of our corrected least squares regression in the 81 unconstrained subjects that are fitted with the $\mathcal{M}^{\text{HARA}}$ model. We explore their behavior as a function of the path as

²⁰The literature usually uses status-quo or lagged status-quo as natural candidates for the reference point. Koszegi and Rabin (2006) model the reference point as an expectation.

well as the returns obtained in the period immediately before.

A. Path Dependence

To test for path dependence, we run the following regression:

$$(7) \quad \hat{u}_{i,t} = \beta_0 + \alpha I_{i,t}^{path>8} + \beta_1 X_{i,t} + \beta_2 F_{i,t} + \rho \hat{u}_{i,t-1} + v_{i,t}$$

where $I_{i,t}^{path>8}$ is a dummy variable that takes a value of 1 if the observation is from a late path (9 to 15) and 0 otherwise. The regression shows no evidence of path dependence for 63 subjects (at the 5% significance level). Among the remaining subjects, 9 exhibit a positive α -parameter, indicating more risk-taking behavior over time than predicted by the model, and 9 exhibit a negative α -parameter, indicating less risk-taking over time than predicted by the model.²¹ A possible explanation is that subjects learn about their preferences over time and adapt their behavior gradually. To investigate this issue further, we run a regression with squared residuals as the dependent variable to assess whether the decisions of subjects become more precise over time. We find that among the 18 subjects with path dependency, one subject commits more mistakes over time (decreasing precision) and no subjects commit fewer mistakes over time. Finally, we examine the α -coefficients of the subjects with a statistically significant effect. The largest positive and negative coefficients are $\alpha = 0.56$ and $\alpha = -0.51$, meaning that the error in the estimation due to path dependency is relatively small. To summarize, 22% of the individuals show statistically significant path dependency, but they go in both directions and are small in magnitude.

²¹The results are similar when we run the regression: $\hat{u}_{i,t} = \beta_0 + \alpha PT_{i,t} + \beta_1 X_{i,t} + \beta_2 F_{i,t} + \rho \hat{u}_{i,t-1} + v_{i,t}$, where the independent variable PT (Path) takes values from 1 to 15.

B. Gain/Loss Asymmetry

To check whether subjects react differently after a loss or a gain, we run the regression:

$$\hat{u}_{i,t} = \beta_0 + \alpha I_{i,t}^{gain} + \beta_1 X_{i,t} + \beta_2 F_{i,t} + \rho \hat{u}_{i,t-1} + v_{i,t}$$

where $I_{i,t}^{gain}$ is a dummy variable that takes a value of 1 if the subject starts the period t after a gain at $t - 1$ and 0 if she starts the period after a loss at $t - 1$ (we use the White-Huber standard errors to account for heteroscedasticity). Our data show no reaction to previous gains or losses beyond the model prediction for 30 subjects (at the 5% significance level). Among the remaining 51 subjects, the vast majority (46 subjects) exhibit higher residuals after a gain. As in Thaler and Johnson (1990), these subjects take more risks after a gain than after a loss. The remaining 5 subjects exhibit the opposite pattern. We then study the magnitude of the α -coefficient for subjects with a significant overreaction to previous outcomes. Among subjects who take more risks after a gain, 40 have a small overreaction (\$1 or less) and 6 have a more substantial one (between \$1 and \$4). All 5 subjects who take more risks after a loss have a small coefficient: $|\alpha| < 0.48$. In summary, while many subjects (57%) exhibit excessive risk-taking after gains, the overreaction is small in magnitude with some exceptions (7% of subjects).

C. Summary of the Behavioral Types

It is interesting to notice that the two sets of anomalies involve mostly different subjects: 6 individuals exhibit path dependence, 39 exhibit a gain/loss asymmetry, and only 12 exhibit both anomalies. Also, subjects with one or both anomalies are present in all of the risk-type categories described in Table 3. The remaining 24 subjects can be very

confidently classified as expected utility maximizers.

In conclusion, anomalies are prevalent. Residual behavior can be attributed to systematic biases that are not captured by the structural model. At the same time, anomalies are spread among subjects and are small in magnitude, so we can fit the data to the expected utility model reasonably well.

VI. Session Level Risk Attitudes

To assess the group level risk attitudes, we now perform the same classification exercise as in section C except that we conduct the analysis at the *session* level rather than at the individual level. Since we are not aware of any established methodology to perform a group level analysis in the setting of dynamic portfolio allocation, we explore two approaches.

The per-capita agent. Recall that our experiment consists of 13 sessions with 7 to 10 participants each. This design permits an objective measure of aggregate wealth because all participants within a session are subject to the same shock. Instead of summing all the wealth accumulated by subjects in each period, we adopt a per-capita specification, which allows us to identify the risk preferences of the “per-capita agent.”²² Accordingly, the per-capita amount invested in the risky asset in each period represents per-capita agent’s

²²The per-capita agent is a description of the group risk attitudes and is not related to any well-known theoretical construct of a representative agent, which is intrinsically a market equilibrium concept. In a recent market experiment, Asparouhova, Bossaerts, Roy, and Zame (2016) find that individual decisions have little explanatory power for market prices. In other words, the preferences of the representative agent may look nothing like the preferences of the individuals they are composed of. Our description of group risk attitudes is more relevant to a wealth manager who may find it useful to know the preferences of each individual client and the dollar-weighted preferences of the entire group.

allocation decision. We then fit our structural model $\mathcal{M}^{\text{HARA}}$ to the transformed data. This approach makes the results of the individual and session analyses comparable.

The fixed effects approach. We pool all the same session individual level data and apply our structural model $\mathcal{M}^{\text{HARA}}$ in equation 5 with added individual fixed effects and repeat the evaluation for all 13 sessions. The resulting coefficients a and b can then be interpreted as some average risk preferences for the session.

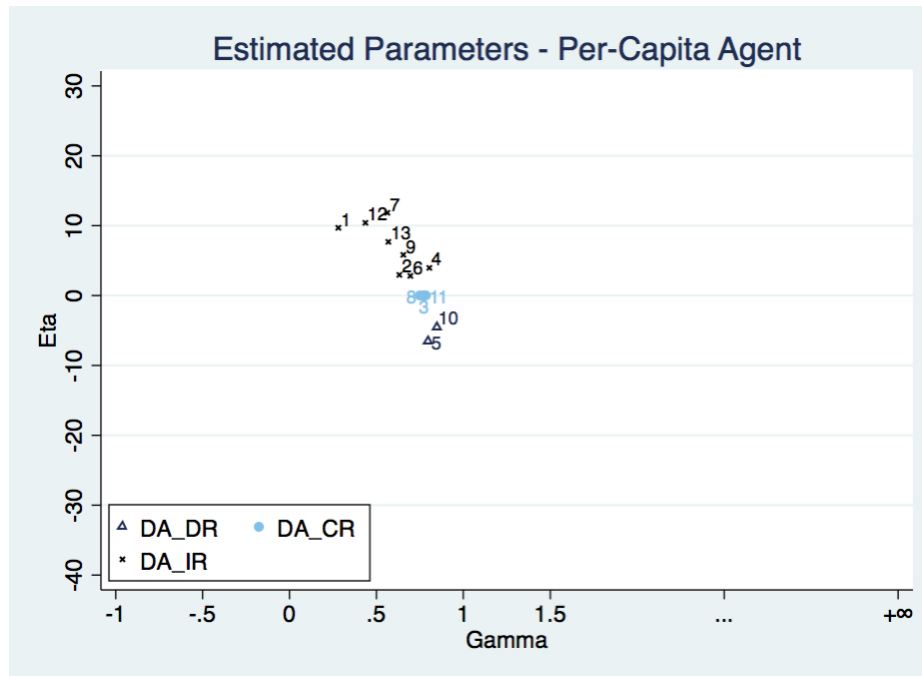
As in section C, we correct for heteroscedasticity and serial correlation using the Newey-West standard errors and estimate γ and η to obtain the risk types at the session level.²³ We conduct this analysis using the data from the 81 unconstrained subjects, so we can draw a comparison between the individual and the session cases.

Figures 3(a) and 3(b) display the analogue of Figure 2 for the session level analysis. The upper panel shows the estimated (γ, η) risk parameters of the 13 per-capita agents of our experiment when the 81 unconstrained subjects are considered. The lower panel shows the estimated (γ, η) risk parameters of the 13 sessions using the fixed effect approach. We label the sessions on the graph to facilitate comparisons of the two approaches. Table 6 reports the analogue of Table 3, that is, the relative and absolute risk aversion at the session level based on the estimated parameters using the per-capita and the fixed effect approach, respectively.

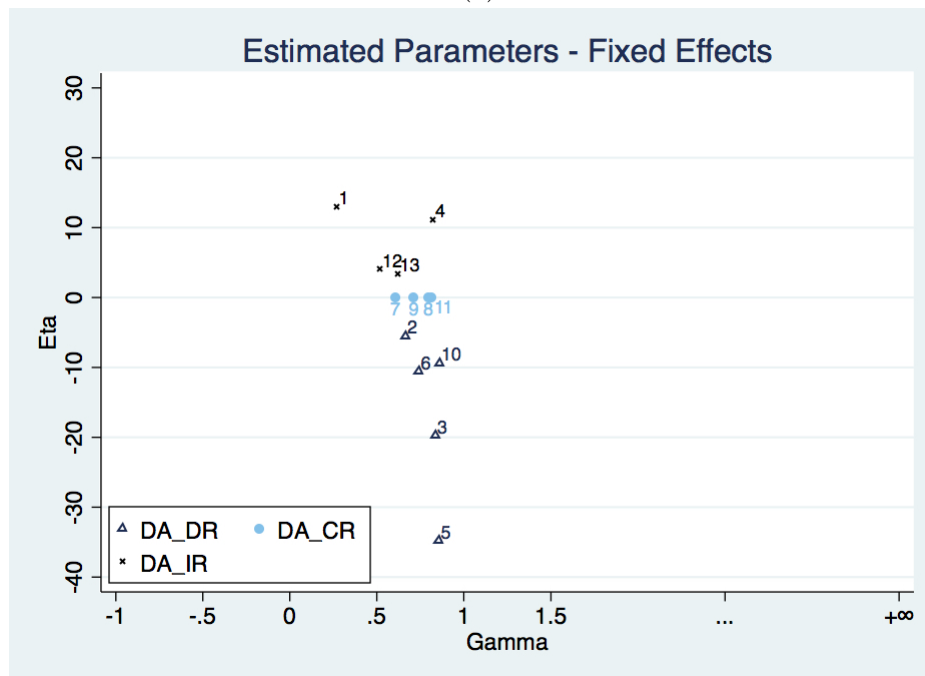
Regardless of the approach, all sessions exhibit DARA. In other words, and unlike the individual level analysis, there is no evidence of CARA at the session level. As for relative risk aversion, we find a reasonably similar proportion of CRRA estimates at the

²³By aggregating wealth, the per-capita investment is always interior. The results of the White test indicate the presence of heteroscedasticity in all sessions. The Breusch-Godfrey test reveals first-order serial correlation in 8 of the 13 sessions.

Figure 3. Estimated Parameters at the Session Level (81 Unconstrained Subjects)



(a)



(b)

Table 6. Session Level Risk Attitudes

Risk attitude	Per-capita agent	Fixed effects
DARA-DRRA	2	5
DARA-CRRA	3	4
DARA-IRRA	8	4
IARA-IRRA	0	0
CARA-IRRA	0	0
# sessions	13	13

individual and session levels (16% to 31%), but significantly more IRRA estimates at the individual or per-capita session (70% and 62%) than at the fixed effect session level (31%).

Even when we only look at the DARA types, the estimates of γ are substantially more dispersed at the individual level ($\gamma \in (-1.2, 0.9)$) than at the session level ($\gamma \in (0.2, 0.9)$) regardless of whether we consider the per-capita or the fixed effects approach.

The estimates of η are, on average, higher at the individual than at the session level.²⁴ This result is important as it suggests that wealth effects are weaker when we aggregate information, either with the per-capita or the fixed effects approach because DRRA agents accumulate, on average, more wealth than IRRA agents and therefore end up having more weight on the session behavior. Therefore, even though there are fewer DRRA than IRRA individuals, their impact in the economy is larger. The effect is exacerbated with the fixed effects approach likely due to the sensitivity of the OLS regression to outlier observations, that is, individuals who accumulate high wealth by investing heavily in the risky asset. Finally, we also estimate the risk parameters of the structural model at the level of the entire experiment with both individual and session fixed effects. We obtain

²⁴The estimates of η are also more concentrated at the per-capita session than at the individual level. The comparison of dispersion is more ambiguous between the fixed effect session and individual level due to the low estimates in sessions 3 and 5.

$\eta^* = -8.7$ and $\gamma^* = 0.76$ suggesting that if we aggregate the entire population, the “average preferences” are best described by DARA-DRRA.

For the per-capita approach, we check for evidence of behavioral anomalies at the aggregate level by replicating the analysis of section V. There is little evidence of path dependence; only three sessions show an effect, and the magnitudes are small. By contrast and consistent with Thaler and Johnson (1990), we find that in all 13 sessions the per-capita agent takes more risk after a gain than after a loss (that is, residuals are higher after a gain). Once again, however, the magnitude of the anomaly is small.

Finally, in Appendix C, we perform the same robustness checks for the per-capita agent as we did for the individual analysis, and obtain similar conclusions. Most notably, the out-of-sample predictions of the HARA model significantly improve those of CRRA in one-third of the sessions and are very similar in the rest.

In summary, while the individual analysis shows that the majority of subjects are DARA-IRRA expected utility maximizers when we perform some type of aggregation, the resulting behavior moves closer to DARA-DRRA or DARA-CRRA types.

VII. Conclusion

In this paper, we report the results of an experiment in which 117 subjects dynamically choose their wealth allocation. Assuming a HARA utility function, we first construct a structural dynamic choice model, which we then use to estimate the absolute and relative risk aversion of the participants. Although technically more complex, this method has the advantage of providing more accurate estimates than traditional risk elicitation techniques.

Even though we find substantial heterogeneity in behavior, decreasing absolute risk aversion and increasing relative risk aversion are the most prevalent subtypes, and we can confidently classify more than half of subjects in the combined DARA-IRRA category. We also find evidence of increased risk taking after a gain but the effect is small in magnitude, and the behavior of subjects is generally well accounted for by the expected utility model. Finally, our design allows us to perform an aggregate analysis. We find that the session level risk attitudes show a different profile than the individual description of risk attitudes, with lower coefficients of relative risk aversion, more concentration in the coefficients of absolute risk aversion, and no evidence of constant absolute risk aversion.

Recent papers have argued that risk attitudes are volatile and difficult to pinpoint (see Friedman et al. (2014) for a survey). Our analysis suggests that if the experimental setting is rich enough, it is possible to accurately estimate (stable) risk preferences. This result is encouraging given the paramount importance for microeconomic theory in understanding risk choices in financial, insurance and environmental settings, to name a few.

References

Abeler, J.; A. Falk; L. Goette; and D. Huffman. “Reference Points and Effort Provision.”

American Economic Review, 101 (2011), 470–492.

Andersen, S.; G. W. Harrison; M. I. Lau; and E. E. Rutström. “Elicitation Using Multiple

Price List Formats.” *Experimental Economics*, 9 (2006), 383–405.

Andersen, S.; G. W. Harrison; M. I. Lau; and E. E. Rutström. “Risk Aversion in Game

Shows.” In *Risk Aversion in Experiments (Vol. 12)*, J. C. Cox and G. W. Harrison, eds.

Bingley, UK: Emerald Group Publishing (2008), 359–404.

Arrow, K. J. *Essays in the Theory of Risk Bearing*. Chicago, IL: Markham (1971).

Asparouhova, E.; P. Bossaerts; N. Roy; and W. Zame. ““Lucas” in the Laboratory.”

Journal of Finance, 71 (2016), 2727–2780.

Assenza, T.; T. Bao; C. Hommes; and D. Massaro. “Experiments on Expectations in

Macroeconomics and Finance.” In *Experiments in Macroeconomics (Research in*

Experimental Economics, Volume 17), John Duffy, ed. Emerald Group Publishing Limited

(2014), 11–70.

Becker, G. M.; M. H. DeGroot; and J. Marschak. “Measuring Utility by a Single-Response

Sequential Method.” *Behavioral Science*, 9 (1964), 226–232.

Binswanger, H. P. “Attitudes Toward Risk: Experimental Measurement in Rural India.”

American Journal of Agricultural Economics, 62 (1980), 395–407.

Bossaerts, P.; P. Ghirardato; S. Guarnaschelli; and W. R. Zame. “Ambiguity in Asset

Markets: Theory and Experiment.” *Review of Financial Studies*, 23 (2010), 1325–1359.

- Bossaerts, P.; D. Meloso; and W. Zame. “Dynamically Complete Experimental Asset Markets.” Working Paper, University of Melbourne (2015).
- Bossaerts, P. and C. Plott. “Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets.” *Review of Finance*, 8 (2004), 135–169.
- Bossaerts, P.; C. Plott; and W. R. Zame. “Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiments.” *Econometrica*, 75 (2007), 993–1038.
- Brocas, I.; J. D. Carrillo; A. Giga; and F. Zapatero. “Skewness Seeking in a Dynamic Portfolio Choice Experiment.” Working Paper, University of Southern California (2017).
- Campbell, J. Y. and J. H. Cochrane. “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior.” *Journal of Political Economy*, 107 (1999), 205–251.
- Charness, G.; U. Gneezy; and A. Imas. “Experimental Methods: Eliciting Risk Preferences.” *Journal of Economic Behavior & Organization*, 87 (2013), 43–51.
- Choi, S.; R. Fisman; D. Gale; and S. Kariv. “Consistency and Heterogeneity of Individual Behavior Under Uncertainty.” *American Economic Review*, 97 (2007), 1921–1938.
- Cox, J. C. and C. F. Huang. “Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process.” *Journal of Economic Theory*, 49 (1989), 33–83.
- Duffie, D. and P. Protter. “From Discrete to Continuous-Time Finance: Weak Convergence of the Financial Gain Process.” *Mathematical Finance*, 2 (1992), 1–15.

- Eckel, C. C. and P. J. Grossman. “Forecasting Risk Attitudes: An Experimental Study Using Actual and Forecast Gamble Choices.” *Journal of Economic Behavior & Organization*, 68 (2008), 1–17.
- Epstein, L. G. and S. E. Zin. “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework.” *Econometrica*, 57 (1989), 937–969.
- Ericson, K. M. M. and A. Fuster. “Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments.” *Quarterly Journal of Economics*, 126 (2011), 1879–1907.
- Friedman, D.; R. M. Isaac; D. James; and S. Sunder. *Risky Curves: On the Empirical Failure of Expected Utility*. Routledge (2014).
- Friedman, D. and R. Oprea. “A Continuous Dilemma.” *American Economic Review*, 102 (2012), 337–363.
- Frydman, C.; N. Barberis; C. Camerer; P. Bossaerts; and A. Rangel. “Using Neural Data to Test a Theory of Investor Behavior: An Application to Realization Utility.” *Journal of Finance*, 69 (2014), 907–946.
- Gill, D. and V. Prowse. “A Structural Analysis of Disappointment Aversion in a Real Effort Competition.” *American Economic Review*, 102 (2012), 469–503.
- Gneezy, U. and J. Potters. “An Experiment on Risk Taking and Evaluation Periods.” *Quarterly Journal of Economics*, 112 (1997), 631–645.

- Harrison, G. W. and E. E. Rutström. “Risk Aversion in the Laboratory.” In *Risk Aversion in Experiments (Vol. 12)*, J. C. Cox and G. W. Harrison, eds. Bingley, UK: Emerald Group Publishing (2008), 41–196.
- Haruvy, E. and C. N. Noussair. “The Effect of Short Selling on Bubbles and Crashes in Experimental Spot Asset Markets.” *Journal of Finance*, 61 (2006), 1119–1157.
- Hey, J. D. and C. Orme. “Investigating Generalizations of Expected Utility Theory Using Experimental Data.” *Econometrica*, 62 (1994), 1291–1326.
- Holt, C. A. and S. K. Laury. “Risk Aversion and Incentive Effects.” *American Economic Review*, 92 (2002), 1644–1655.
- Imas, A. “The Realization Effect: Risk-Taking After Realized Versus Paper Losses.” *American Economic Review*, 106 (2016), 2086–2109.
- Jacobson, S. and R. Petrie. “Learning From Mistakes: What Do Inconsistent Choices Over Risk Tell Us?” *Journal of Risk and Uncertainty*, 38 (2009), 143–158.
- Kahneman, D. and A. Tversky. “Prospect Theory: An Analysis of Decision Under Risk.” *Econometrica*, 47 (1979), 263–291.
- Karatzas, I.; J. P. Lehoczky; and S. E. Shreve. “Optimal Portfolio and Consumption Decisions for a ”Small Investor” on a Finite Horizon.” *SIAM Journal on Control and Optimization*, 25 (1987), 1557–1586.
- Knetsch, J. L. and W.-K. Wong. “The Endowment Effect and the Reference State: Evidence and Manipulations.” *Journal of Economic Behavior & Organization*, 71 (2009), 407–413.

- Kőszegi, B. and M. Rabin. “A Model of Reference-Dependent Preferences.” *Quarterly Journal of Economics*, 121 (2006), 1133–1165.
- Kroll, Y. and H. Levy. “Further Tests of The Separation Theorem and the Capital Asset Pricing Model.” *American Economic Review*, 82 (1992), 664–670.
- Kroll, Y.; H. Levy; and A. Rapoport. “Experimental Tests of the Separation Theorem and the Capital Asset Pricing Model.” *American Economic Review*, 78 (1988), 500–519.
- Levy, H. “Absolute and Relative Risk Aversion: An Experimental Study.” *Journal of Risk and Uncertainty*, 8 (1994), 289–307.
- Levy, H. and M. Levy. “Prospect Theory and the Investment Horizon.” Working Paper, Hebrew University of Jerusalem (2017).
- Lusardi, A.; A. Samek; A. Kapteyn; L. Glinert; A. Hung; and A. Heinberg. “Visual Tools and Narratives: New Ways to Improve Financial Literacy.” *Journal of Pension Economics & Finance*, 16 (2017), 297–323.
- Maier, J. and M. Rüger. “Measuring Risk Aversion Model-Independently.” Working Paper, University of Munich (2012).
- Merton, R. C. “Optimum Consumption and Portfolio Rules in a Continuous-Time Model.” *Journal of Economic Theory*, 3 (1971), 373–413.
- Plott, C. R. and S. Sunder. “Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets.” *Econometrica*, 56 (1988), 1085–1118.

- Post, T.; M. J. Van den Assem; G. Baltussen; and R. H. Thaler. “Deal or No Deal? Decision Making Under Risk in a Large-Payoff Game Show.” *American Economic Review*, 98 (2008), 38–71.
- Rapoport, A. “Effects of Wealth on Portfolios Under Various Investment Conditions.” *Acta Psychologica*, 55 (1984), 31–51.
- Rapoport, A.; R. Zwick; and S. G. Funk. “Selection of Portfolios With Risky and Riskless Assets: Experimental Tests of Two Expected Utility Models.” *Journal of Economic Psychology*, 9 (1988), 169–194.
- Smith, V. L.; G. L. Suchanek; and A. W. Williams. “Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets.” *Econometrica*, 56 (1988), 1119–1151.
- Sokol-Hessner, P.; M. Hsu; N. G. Curley; M. R. Delgado; C. F. Camerer; and E. A. Phelps. “Thinking Like a Trader Selectively Reduces Individuals’ Loss Aversion.” *Proceedings of the National Academy of Sciences*, 106 (2009), 5035–5040.
- Sundali, J. A. and F. Guerrero. “Managing a 401 (k) Account: An Experiment on Asset Allocation.” *Journal of Behavioral Finance*, 10 (2009), 108–124.
- Thaler, R. H. and E. J. Johnson. “Gambling With the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice.” *Management Science*, 36 (1990), 643–660.
- Wang, S. W.; M. Filiba; and C. F. Camerer. “Dynamically Optimized Sequential Experimentation (DOSE) for Estimating Economic Preference Parameters.” Working Paper, California Institute of Technology (2010).

Appendix

Appendix A. Instructions

Note: The following instructions are accompanied by a slideshow presentation. Slides available upon request.

We are about to begin. Please put your cell phones and other electronic devices in your bag and do not use them until the end of the experiment.

Dear Participants,

Welcome and thank you for coming to this experiment. You will be paid for your participation, in cash, at the end of the experiment. You will remain anonymous to me and to all the other participants during the entire experiment; the only person who will know your identity is the person in the other room who is responsible for paying you at the end. Everyone will be paid in private and you are under no obligation to tell others how much you earned. The entire experiment will take place through the computer terminals.

Let us begin with a brief instruction where you will be given the complete description of the experiment and shown how to use the software. Please, pay attention to the instructions, as it is important for you to understand the details of the experiment. There will be a quiz at the end of the instructions that everyone needs to answer correctly before we can proceed to the actual experiment. Participants who are unable to answer the quiz will not be allowed to participate in the experiment. If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear. If any difficulties arise after the experiment has begun, raise your hand, and an experimenter will come and assist you. If you cannot see the entire projection screen, please come forward as it is important for you to see the entire screen.

Today, we will ask you to make investment decisions. Your final payment consists of a \$5 show-up fee plus your investment earnings. Those earnings depend both on the choices you make

and on luck. The choices of other participants do not affect your payoff in ANY way, at ANY point in the experiment and your choices do not affect their payoff. The entire experiment is split in 3 parts. I will now give you instructions for Part 1. You will get additional instructions before Part 2 and Part 3.

PART 1

Let me first summarize the investment process and then we will go through each step in more detail. You will start with an initial amount of money that you will be able to invest in two assets, A and B. You will have 10 periods to invest. At each period you will decide how to allocate your money between the two assets. At the end of each period, you will earn returns from that period's investment in each asset. The two assets will pay differently, and later in the instructions I will explain what to expect from each asset. After period 10, the process ends and the computer will record your final money amount. This process of 10 investment periods is called an investment path. At the start of each path, your money will be reset to the initial amount.

In Part 1 you will complete 15 of these paths. Consequently, there will be 15 final amounts of money, one for each path. The computer will randomly select one of these 15 final amounts. The selected amount will be your payoff for Part 1. Are there any questions? Let me now walk you through the procedure step by step.

The Initial Endowment: This is a screenshot of what you will see on your computer at the beginning of each path, that is, in period 1 of each path. In each path you start with \$3, this is your initial endowment. This amount is displayed on the left side of the screen in the box labeled "Account". It is also represented by the height of the bar in the middle of the screen. There is a grid in the background to help you get a sense of the bar's height.

Periods and Timing: As mentioned earlier a path is made of 10 periods, starting at 1 and ending at 10. The sequence is displayed on the bottom of the screen and your current period is displayed in the upper left corner. Each period is an opportunity for you to invest. A period ends

when the time runs out. You can see the timer on the left hand side. For the first period in each path you will have 10 seconds to make your investment decision. For the other periods in the path, that is, periods 2 to 10, you will have 6 seconds to make your investment decision.

Investing: Let me show you with a short video how to make your investment.

Step 1. Choose display. To start your investment, you first need to click on one of the two boxes at the bottom, the ones labeled with percentage and dollar signs. These boxes control how your allocation between assets A and B is displayed: in percentages or dollar terms. You have to click one of the boxes in period 1 of each path. You can also change the display anytime simply by clicking on the other box. Select whichever box you find convenient and change it anytime you want.

Step 2. Activate the bar. Now you can activate the investment bar. Click anywhere on the light gray bar to activate it. Notice when the bar is light gray it means that your money is not invested in either asset. If the bar stays that way after the period ends, you will earn zero interest on your money: the same amount will just be transferred to the next period.

Step 3. Choose the allocation. Once you activate the bar you will notice that it is split between two colors: the top is light blue, and the bottom is gray. The top represents the amount of money allocated to Asset A, and the bottom represents the amount of money allocated to Asset B. Now you can see the display I previously mentioned. It shows how much money you have allocated to A and B either in dollar or in percentage terms. This example shows the dollar display. You can change the allocation between A and B in two ways: by holding the horizontal bar and moving it up or down or by clicking on the bar, as you can see in the video. Once you are satisfied with the allocation wait until the period ends.

Step 4. Proceed to the next period. When the period ends, a new gray bar will appear showing you the new amount. Here is the transition from period 1 to period 2. Your new amount will be the sum of the money you earned on both assets A and B and it will be shown on the left

where your initial money amount was displayed. The new height of the bar will also represent this amount. Be aware that the background grid can be re-scaled to accommodate changes in the bar, so pay attention to the figures written on the grid. The last period's bar will become inactive but you will still be able to see your past allocation between assets A and B. Remember that you need to activate the bar and choose an allocation between A and B at every period, otherwise you earn no interest. Here is a period 2 allocation process and the transition to period 3. Notice how I changed the display from dollars to percentages. This process continues until period 10. After period 10, the path ends. Here is a screenshot of one path end. Your final amount will be shown in the box on the left and by the height of a green bar on the right. A message will appear informing you that the path ended. You need to click the "OK" button to continue. A new path will start shortly thereafter.

Assets A and B: Let me show you what to expect from the investment in each asset. In the upper left corner of your screen there is a box that reads "Asset A: mean return 6%, standard deviation 55%"; "Asset B mean return 3%". These numbers show how your investment in each asset grows and they will not change during the entire experiment.

Asset B: The 3% next to Asset B in the box means that, once the period ends, the amount allocated to Asset B will grow by 3% for sure. The interest rate of 3% will not change throughout the duration of the experiment. A reminder: money in Asset B is represented by the bottom, GRAY portion of your active bar. Here is an example: if you have 2 dollars invested in B you will have 2 dollars and 6 cents in the next period. If you keep that money in B you will then have 2 dollars and 12.2 cents the period after. You can think of your money in Asset B being multiplied by 1.03. Note that 2 dollars is just an example. In the experiment you can choose any allocation you want provided it does not exceed your total amount.

Asset A: Contrary to asset B, your return on asset A is uncertain. Technically, the return on asset A has a Normal distribution with mean 6% and standard deviation 55%, as shown in the

upper left box. This means that asset A grows by 6% on average. However, it may be more or it may be less. In particular, the growth rate could be negative. In this case the money you invested in Asset A will shrink. Although the return can be negative, the amount of money you hold on asset A can never go below zero. A reminder: money in Asset A is represented by the top, LIGHT BLUE portion of your active bar.

Another way to think about the return on this asset is that the amount you put in asset A will be multiplied by some positive number. On average, this number will be 1.06 which corresponds to a 6% growth. Let us call this number a multiplier. If the multiplier is less than 1, it means that your investment in Asset A shrinks. For example, if you allocate \$2 to asset A and the multiplier turns out to be 0.8, you will have \$1.6 in the next period. If the multiplier turns out to be 1.5, you will have \$3 the next period. Here is a chart showing the probability of your multiplier being in a given range. With 20% chance it will be somewhere between 0 and 0.67. With 30% chance it will be somewhere between 0.67 and 1.06. With another 30% chance it will be somewhere between 1.06 and 1.7. Finally, with 20% chance it will be above 1.7. Once the period ends and you receive the returns on your assets, the box on the left marked “Last Period Multiplier” will show what turned out to be the multiplier for asset A in that period. The box will show always 1.03 as the multiplier for asset B.

Projection Bar: The returns from asset A obtained after several periods depend on many factors. In order to help you get an idea of the range of outcomes, we placed a projection bar at the end of the screen. Let me explain how the projection bar works. Suppose for example that in the first period you invest \$2 in Asset A. If you keep the returns on that same asset, how much money will you have at the end of the 10th period? Observe what happens on the left hand side of the graph. It is a simulation of your return. The vertical axis represents dollars and the horizontal axis the periods. It begins with 2 dollars in the first period and it ends after 10 periods. Here is one potential final amount of money. But it can also be this. Or this. Or this. Notice that each

time a path ends, we keep track where it lands by adding a dot on the right graph. Each dot represents a possible final amount of money. If we run enough paths, all with \$2 invested in asset A, we will get a bunch of dots on the right end. The more dots each dollar region has, the more likely your amount of money will end up there. And that's exactly what the bar represents: the likelihood of your earnings ending up in a certain amount.

Now look at the example in the picture. It is period 4. Look at the projection bar. For the current investment strategy, the lower gray part is the projection of how much you will earn on asset B if you don't change the allocation between assets until the end of the path. In this case, you will earn 1.89 dollars on asset B. This amount is for sure since there is no uncertainty on this asset. The upper part shows the projected earnings on asset A if you don't change the allocation between assets until the end of the path. They correspond to the dots shown in the video. There is a 20% chance that the final amount lands in the white area above the gray one, a 60% chance that it lands in the dark blue area and another 20% chance that it lands above the dark blue area. Finally, there is a thick line showing the median, in this example, 16 dollars and 64 cents. This means that with a 50% chance your final amount will be somewhere below that number and with a 50% chance it will be somewhere above that number.

Notice also from our demonstration that probabilities are different within a segment. For example, receiving an amount above the dark blue area has a 20% chance, but within this 20% it is more likely to be close to the dark blue area than further away. In other words, it is more likely to get this payoff [point to the slide] than this payoff [point to the slide], although both are possible. You can see this point more clearly on the frequency table. Based on the number of circles, it is more likely that your payoff will end up here [point to the slide] rather than [point to the slide] here, even though both of these areas correspond to the 20% region above the projection bar.

Important Points:

1. The projection bar shows the likelihood of different final earnings at the end of the path

ASSUMING the amount you receive from each asset is reinvested in the same asset in all the following periods. However, you can change the allocation between assets at every period.

2. At each period, the projection bar recalculates the probabilities. If you move the cursor up and down within a period, the bar shows instantly the new projection.
3. The multipliers on asset A are independent across periods. In other words, the multipliers of previous periods will in no way impact the multiplier in the current period. For example, if the multiplier in a previous period was very high, it does not mean it will be high again. The new multiplier will simply follow the rules of uncertainty described before.
4. All the participants start and end the paths at the same time. The clock starts as soon as the screen appears, so pay attention.
5. The multiplier for asset A in each period is the same for all participants. So, for example if the computer chooses 2 as a multiplier in period 4, it means that all participants will have their investment in asset A doubled.

Are there any questions? If not, let us proceed to 5 practice paths. What you earn on these paths will not count towards your payment; these are meant only for you to familiarize yourself with the entire process of allocating money between assets A and B. Feel free to explore as many investment strategies as possible to better understand the different options.

Please double click on the icon on your desktop that says ABC STUDY. When the computer prompts you for your name, type a 4 digit number that you can easily remember. Please do not forget the number you typed. Then click SUBMIT and wait for further instructions.

Pay attention to the screen. The first practice path will be starting soon. Focus on understanding how to choose the display between percentage and dollars, how to activate the bar, and how to change the allocation between assets. Reminder: Once a path ends, you need to click the OK button in order to proceed to the next path.

[START game] [Complete practice path 1]

You have now completed practice path 1. Are there any questions? Let's proceed to practice paths 2 and 3. Now try to explore different investment strategies to get a good understanding of the investment process.

[Complete practice paths 2 and 3]

You have now completed 3 practice paths. Are there any questions? If not, we will proceed to a short quiz. Please pay close attention to answering the questions, as you will not be permitted to continue with the experiment if you do not answer the questions correctly. Raise your hand if you have any question during the quiz.

[Complete quiz]

You have now completed the quiz. Let us proceed to the last 2 practice paths.

[Complete practice path 4 and 5]

You have now completed practice paths 4 and 5. Are there any questions? Before we start please write down your ID on your record sheet in front of you. You will locate your ID on the left side of your window bar. You will have to present the record sheet to get paid at the end of the experiment. Did everyone right their IDs down?

Let me remind you how you will be paid for Part 1. At the end of the experiment, the computer will randomly select one of the 15 paths and you'll be paid the final amount you earned in that path. Are there any questions? If there are any problems or questions from this point on, raise your hand and an experimenter will come and assist you. We are ready to start the experiment. Please pull out your dividers.

QUIZ (accompanied by a display print-out):

1. Look at the display on the paper in front of you. What is the current period?

(a) 1

(b) 2

(c) 6 (correct)

(d) 8

2. Had the person not chosen any allocation between assets A and B, how much would she have in the next period?

(a) \$1.44

(b) \$4.16

(c) \$5.60 (correct)

(d) \$7.29

3. In this period, how much has the person invested in Asset A?

(a) \$5.60

(b) \$4.16 (correct)

(c) \$1.44

(d) \$0.81

4. Assume this person keeps reinvesting the returns of asset A in A and the returns of asset B in B until the end of the path. Given the current allocation, how much money will this person have in asset B after the path ends?

(a) \$1.67 (correct)

(b) \$1.44

(c) \$1.03

(d) \$7.29

5. Assume this person keeps reinvesting the returns of asset A on A and the returns of asset B on B until the end of the path. Given the current allocation, how much money will this person have in asset A after the path ends?
- (a) \$7.29
 - (b) \$5.60
 - (c) \$18.00
 - (d) Cannot be determined with certainty (correct)
6. Forget about the display. Imagine you invest \$1 in Asset A and \$1 in Asset B and suppose the multiplier on Assets A and B are 2.00 and 1.03 respectively. How much money will you have in the next period?
- (a) \$3.03 (correct)
 - (b) \$4.00
 - (c) \$5.03
 - (d) Cannot be determined from the information given.

Appendix B. Individual Analysis: Robustness and Extensions

In this appendix, we explore the validity of the HARA specification, both overall and in comparison to the one-parameter CRRA specification. We restrict our attention to the 81 subjects whose behavior can be fitted to our structural $\mathcal{M}^{\text{HARA}}$ model.

To assess the general goodness of fit of our structural model, we first conduct F-tests to determine whether the proposed structural relationship between the risky investment at each period and the set of independent variables is statistically reliable. We find that it is for all 81 subjects. Furthermore, according to the Akaike Information Criterion comparison (AIC), $\mathcal{M}^{\text{HARA}}$ outperforms $\mathcal{M}^{\text{CRRA}}$ for all but 7 subjects.²⁵

B.1 Out-of-Sample Predictions and Comparison HARA vs. CRRA

We study whether we can predict the behavior of our subjects based on the observation of their choices in a subset of trials. More specifically, we randomly choose eight paths to estimate the parameters of the $\mathcal{M}^{\text{HARA}}$ model and then use the estimates to predict choices on the remaining seven paths. We repeat this exercise 100 times. For each repetition, we calculate the Mean Absolute Error:²⁶

$$MAE^{\text{HARA}} = \frac{\sum_{i=1}^7 \sum_{t=1}^{10} |\pi_{i,t} - \tilde{\pi}_{i,t}^{\text{HARA}}|}{70}$$

where $\tilde{\pi}_{i,t}^{\text{HARA}}$ is the prediction of $\mathcal{M}^{\text{HARA}}$ on decisions in the 7 validation paths.

²⁵The η parameter of $\mathcal{M}^{\text{HARA}}$ is estimated to be zero for these 7 subjects, implying de-facto constant relative risk aversion. If we use the Bayesian Information Criterion (BIC), there are 4 more subjects for which $\mathcal{M}^{\text{CRRA}}$ outperforms $\mathcal{M}^{\text{HARA}}$. The estimated η parameter is zero or close to zero for all 11 subjects.

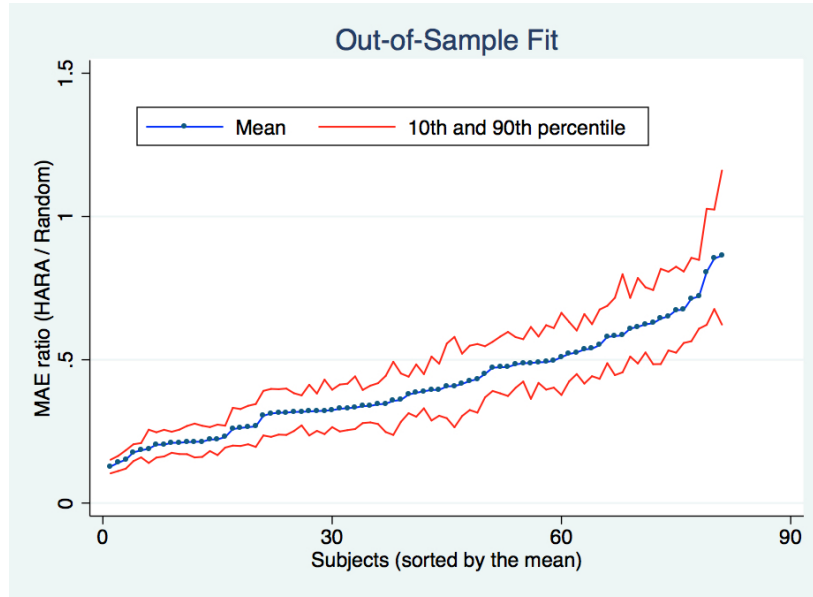
²⁶We conducted the same analysis with the root mean square error measure (RMSE) instead of the MAE and obtained similar results.

As a benchmark, we first compare the out-of-sample fit of HARA to a model where fit decisions are randomly made for the same seven validation paths. More specifically, we calculate:

$$MAE^{\text{RND}} = \frac{\sum_{i=1}^7 \sum_{t=1}^{10} |\pi_{i,t} - \tilde{\pi}_{i,t}^{\text{RND}}|}{70}$$

where $\tilde{\pi}_{i,t}^{\text{RND}}$ is an amount drawn from a uniform distribution in $[0, X_{i,t}]$. Figure 4 shows the mean, 10th percentile, and 90th percentile of the ratio $\frac{MAE^{\text{HARA}}}{MAE^{\text{RND}}}$ of 100 repetitions for our 81 subjects, sorted by the mean of the ratios, from smallest to largest. The ratio is intended to describe how much better $\mathcal{M}^{\text{HARA}}$ explains behavior in comparison to a naïve random model. The ratio is below 1 for all subjects and below 0.5 for 73% of subjects, suggesting (not surprisingly) that for the vast majority the HARA specification performs substantially better out-of-sample than the random specification.

Figure 4. Out-of-Sample Fit (100 Repetitions) - HARA vs. Random



Next, we ask a more relevant question: how much predictive power do we lose by considering a simple, one-parameter CRRA specification instead of the richer, two-parameter HARA specification? We follow the same procedure as before with the

$\mathcal{M}^{\text{CRRA}}$ model and calculate the Mean Absolute Error:

$$MAE^{\text{CRRA}} = \frac{\sum_{i=1}^7 \sum_{t=1}^{10} |\pi_{i,t} - \tilde{\pi}_{i,t}^{\text{CRRA}}|}{70}$$

where $\tilde{\pi}_{i,t}^{\text{CRRA}}$ is the prediction of $\mathcal{M}^{\text{CRRA}}$ on decisions in the 7 validation paths. Figure 5 shows the mean, 10th percentile, and 90th percentile of the ratio $\frac{MAE^{\text{HARA}}}{MAE^{\text{CRRA}}}$ of 100 repetitions for our 81 subjects.

Figure 5. Out-of-Sample Fit (100 Repetitions) - HARA vs. CRRA



The mean of ratios is smaller than 1 for 61 subjects (75%). Half of these subjects have mean ratios below 0.9 and at least 90% of repetitions below 1, which suggests that the improvement of HARA over CRRA is substantial for 37% of subjects and minor for the other 38%. Among the remaining 20 subjects for whom CRRA performs better out-of-sample (25% of the population), the mean ratio is above 1.1 for only 1 subject. Notice also that 11 out of those 20 subjects are specified as having constant relative risk aversion by the HARA model ($\eta = 0$), so the similarity between the out-of-sample predictions of the two models is expected for those individuals. Overall, HARA improves

significantly out-of-sample predictions over CRRA for one-third of the sample and performs similarly for the other two-thirds.

B.2 Risk Type Predictions: Early vs. Late Paths

Our second prediction exercise consists of determining whether the risk type obtained from the data in one sample is consistent with the risk type obtained in the complement. If they are not, it might be because of learning or a preference change. To do so, we divide our sample into “early paths” (first 8 paths) and “late paths” (last 7 paths). We then estimate the risk types of the individuals in each subsample following the same methodology as before, and look at the consistency in the classification of subjects across datasets. The results are reported in Table 7.

Table 7. Consistency in Classification - Early Paths vs. Late Paths

Type consistency by paths		Frequency
Consistent	Full - Early - Late	46
Consistent	Full - Early	21
Consistent	Full - Late	11
Inconsistent		3
Total		81

In line with previous results, we find that subjects are fairly consistent between early and late paths. Of the 81 subjects, 46 have the same risk type across all samples, 32 are consistent on the full sample and one subsample, and only 3 subjects have different risk types in all samples. Table 8 presents the risk aversion attitude of subjects in the full sample as well as the early path and late paths subsamples.

Table 8. Type frequency - Early Paths vs. Late Paths

	Full sample	Early paths	Late paths
DARA-DRRA	11	12	10
DARA-CRRA	13	14	23
DARA-IRRA	44	38	33
IARA-IRRA	1	1	1
CARA-IRRA	12	16	14

The proportions of the different risk types are, to a large extent, preserved in all

samples: there is a majority of DARA-IRRA (between 41% and 54%), virtually no IARA-IRRA, and some representation of the other three types. The most notable difference between samples is the increase in CRRA types at the expense of IRRA types. However, we do not want to excessively emphasize this conclusion as it may be partly due to a statistical effect: with fewer observations per subject in each subsample it may be more difficult to reject the null hypothesis of constant relative risk aversion.

Overall, the expected utility model performs well. Subjects do not change risk types dramatically over the course of the experiment and it is possible to predict with reasonable accuracy their behavior after observing the choices in the first paths.

B.3 Risk Type Predictions: Early vs. Late Periods

We perform the same analysis as in appendix B.2, except that we divide the sample into early periods (first 5 of each path) and late periods (last 5 of each path). Notice that wealth is likely to be lower in early periods. We first confirm this intuition: wealth is, on average, \$6.7 in the last five periods compared to \$3.9 in the first five periods for the 81 unconstrained subjects. With this potential source of differences in mind, we present in Tables 9 and 10 the analogue information of Tables 7 and 8 for the new subdivision of samples.²⁷

As before, most subjects are consistent between the full sample and at least one subsample, though the number of inconsistent subjects is slightly higher than for the early/late path division (8 vs. 3). Risk attitudes are also similar between samples, with a large representation of DARA-IRRA and an absence of IARA-IRRA. Also, as before (and

²⁷The type “Not classified” in Table 10 corresponds to individuals for which both the a and b coefficients in equation (5) are zero.

Table 9. Consistency in Classification - Early Periods vs. Late Periods

Type consistency by periods		Frequency
Consistent	Full - Early - Late	39
Consistent	Full - Early	9
Consistent	Full - Late	25
Inconsistent		8
Total		81

Table 10. Type Frequency - Early Periods vs. Late Periods

	Full sample	Early periods	Late periods
DARA-DRRA	11	11	8
DARA-CRRA	13	27	20
DARA-IRRA	44	31	37
IARA-IRRA	1	1	1
CARA-IRRA	12	10	14
Not classified	0	1	1

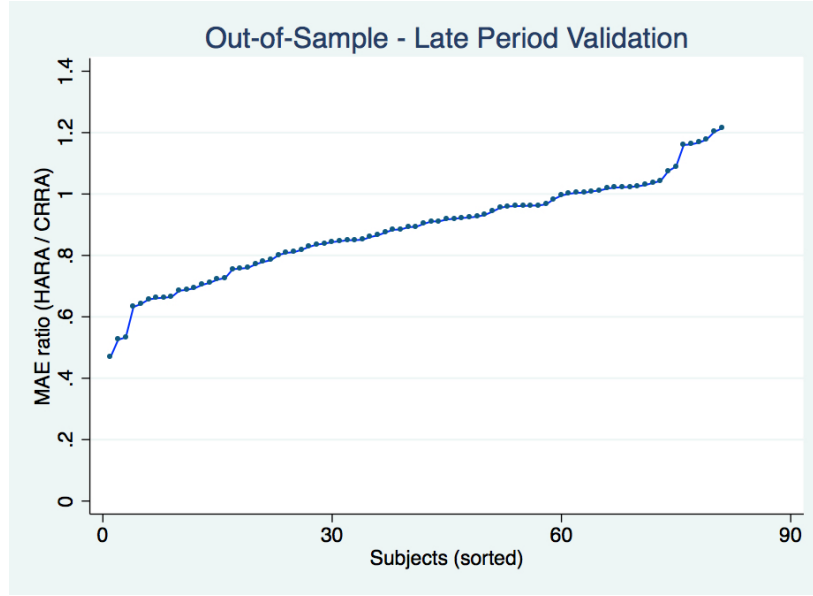
with the same caveat), there is an increase in CRRA subjects at the expense of IRRA subjects.

B.4 Out-of-Sample Predictions: Low vs. High Wealth

We conduct a similar out-of-sample prediction analysis as in appendix B.1, except that we estimate parameters on the first five periods of all paths and predict choices in the last five periods of all paths. This is different than before not only in that the sample division is based on periods rather than paths but also in that we do not randomly take subsamples. The purpose is to test whether extrapolating choices based on risk attitudes elicited in trials with low wealth is meaningful to explain decisions in trials with high wealth.

We find that all but one subject has a ratio $\frac{MAE^{HARA}}{MAE^{RND}}$ smaller than 1 and, as before, more than two-thirds of subjects have a ratio smaller than 0.5, indicating (not surprisingly) a large improvement of HARA over random choice. We then present in Figure 6 the analogue of Figure 5 to the new sample division, that is, the the ratio $\frac{MAE^{HARA}}{MAE^{CRRA}}$ sorted by subjects from smallest to largest.

Figure 6. Out-of-Sample Fit - HARA vs. CRRA in Early vs. Late Periods



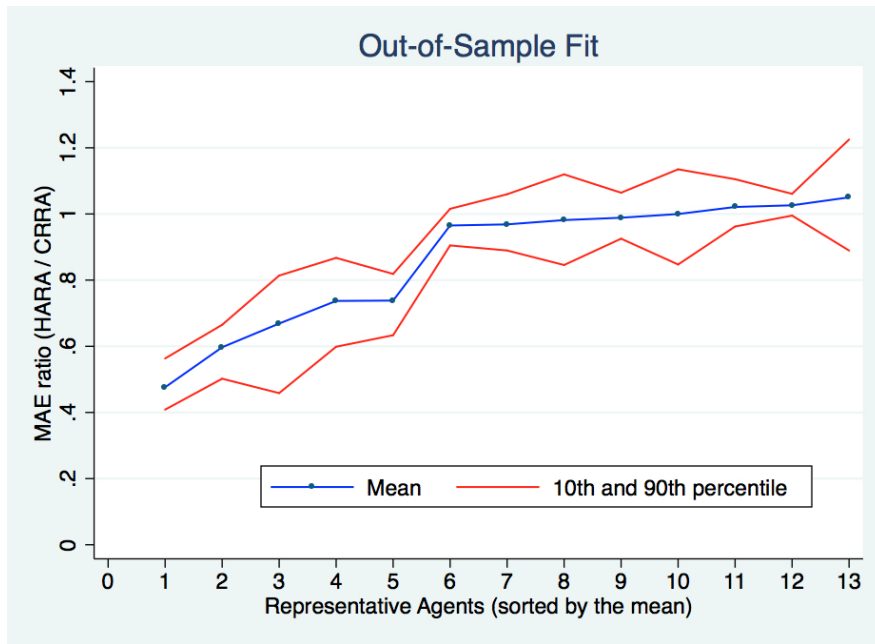
The results are remarkably similar to those obtained in appendix B.1. One-third of subjects exhibit a considerable improvement of HARA over CRRA whereas the other two-thirds are similar. The most notable difference is the existence of a 6 subjects for which CRRA performs better than HARA. Overall, the results confirm those above; the estimated types are consistent across subsamples (even when we use “low” wealth estimates to predict “high” wealth choices) and the general utility function helps in the estimation for one-third of the individuals.

Appendix C. Session Level Analysis: Robustness and Extensions

We explore the out-of-sample predictive properties of the model by performing the same analysis as in appendix B.1. Again, as benchmark, we compare HARA to random choice. For all 13 per-capita agents, the ratio $\frac{MAE^{HARA}}{MAE^{RND}}$ is below 0.5 in at least 90% of the repetitions. This means that the improvement of HARA over random choice is greater for the per-capita agent than for the individual analysis, which implicitly suggests that some of the subjects’ deviations cancel each other out.

We then analyze how HARA compares to CRRA. Figure 7 is analogous to Figure 5 and shows the mean, 10th percentile, and 90th percentile of the ratio $\frac{MAE^{HARA}}{MAE^{CRRA}}$ of 100 repetitions for the 13 per-capita agents. The mean ratio is virtually 1 for 8 sessions and between 0.5 and 0.8 for the other 5 sessions. This means that, as for the individual analysis, the out-of-sample predictions of the HARA model significantly improve those of CRRA for one-third of the sample and are very similar for the rest.²⁸

Figure 7. Out-of-Sample Fit (100 Repetitions) - HARA vs. CRRA



When we divide the sample between early paths and late paths, as we did in appendix B.2 for the individuals, we find that the representative agent has the same risk type across all samples in 7 sessions (4 DARA-IRRA and 3 DARA-CRRA). Of the remaining 6 sessions, 2 sessions show the same type in the full and early paths subsample and 4 sessions show the same type in the full and late paths subsample. Overall, representative agents are generally consistent across paths, a result that is not surprising

²⁸When we include the constrained subjects, we find a larger improvement of HARA over CRRA. This is due mostly to the fact that the vast majority of the constrained subjects exhibit IRRA behavior.

given the type-consistency of the majority of individuals in our sample.