Identifying Time Preferences With Experiments: Comment

by

Glenn W. Harrison, Morten I. Lau & E. Elisabet Rutström†

February 2013

ABSTRACT.

Identifying time preferences with laboratory experiments demands attention to theoretical, experimental and econometrics issues. Andreoni and Sprenger [2012a] propose a single choice task and several econometric methods that seek to address these issues. The choice task requires subjects to make portfolio allocations between sooner and later payments of money. All theories they examine imply that subjects pick one boundary or the other, or that they pick strictly interior allocations. Their econometric methods seek to explain the average portfolio choice, but ignore the bald fact that 70% of the responses by the subjects were choices at one boundary allocation or the other. The average portfolio choice implied by the modes at either boundary is chosen by virtually none of their subjects. Their *ad hoc* econometric attempts to model the truncation of choices at the boundaries fail to account for the economics of the observed behavior. A systematic analysis of their data generates *a priori* implausible estimates of significantly convex utility functions. Andreoni and Sprenger [2012b] inherits all of the problems of the basic design and econometrics from Andreoni and Sprenger [2012a], and adds one: their findings are immediately confounded by non-additivity of the intertemporal utility function. Apart from this theoretical confound, there is experimental evidence of just this type of non-additivity, leading to an aversion to correlated payoffs over time. The evidence in favor of correlation aversion predicts the qualitative pattern they observe perfectly, without claiming that the utility function for stochastic outcomes is somehow different from the utility function for non-stochastic outcomes.

JEL Codes: D91, C91, D81

† Department of Risk Management & Insurance and Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA (Harrison); Durham University Business School (Lau); and Robinson College of Business, Georgia State University, USA (Rutström). E-mail contacts: gharrison@gsu.edu, m.i.lau@durham.ac.uk and erutstrom@gmail.com.
We thank the Danish Social Science Research Council (projects 275-08-0289 and 12-130950) for research support.

Table of Contents

The identification of time preferences with controlled experimental tasks involves a careful interplay of theory, experimental design, and econometric inference. While experimental data on risk and time preferences has greatly moved our understanding of these preferences forward, there is still a need for a careful and productive debate about methods and the claims they have generated. We contribute to this debate by critically evaluating the methods and findings of Andreoni and Sprenger (A&S) [2012a][2012b].

There are two popular, competing approaches to eliciting time preferences using monetary incentives. The main differences are that one approach uses binary choices of two types of outcomes and the other approach uses portfolio choices over one type of outcome. We refer to the former as the Binary Choice (BC) and the latter as the Convex Time Budget (CTB) approach. Early uses of BC with monetary incentives are Coller and Williams [1999] and Harrison, Lau and Williams [2002]. Subjects are given binary choices between sooner, smaller amounts of money and larger, later amounts of money. In the simplest designs, the sooner amount is kept constant and the later amounts increase in an ordered manner. For present purposes, there are two key features of this design: the choices are binary, and the payoffs are non-stochastic. This approach greatly reduced the very high discount rates that had been elicited in experiments with small or hypothetical stakes and open-ended elicitation methods: from hundreds or thousands of percent per annum down to 25% or 30% p.a.. The non-stochastic approach was later extended by Andersen, Harrison, Lau and Rutstrom (AHLR) [2008a] to include binary choices over lotteries that differ in risk but include no time delay. This extension allowed risk and time preferences to be jointly estimated, correcting the inferred discount rates for the otherwise confounding presence of utility concavity. The extended approach lead to even further reductions in elicited discount rates, placing these firmly in a range that is considered reasonable: around 10% p.a..

The CTB approach for eliciting time preferences was introduced by A&S [2012a]. It allows subjects in each choice to choose a portfolio of sooner payoffs and later payoffs, all non-stochastic. A&S

[2012b] extend this approach to include stochastic payoffs over time. The realizations in the stochastic design occurred at the same time, so that no preferences for temporal resolution of risk would confound inferences (Kreps and Porteus [1978]). The key innovation here over the earlier designs follows Cubitt and Read [2007; p.387], and allows for portfolio allocations within the two endpoints of each binary choice. While acknowledging the theoretical logic behind the AHLR [2008a] design that controls for utility concavity, they include no additional task and base their utility corrections on parameter estimates from the choices over sooner and later payoffs.

We criticize here three claims in A&S [2012a][2012b]. The first claim is that Nonlinear Least Squares or Tobit are appropriate econometric techniques to use on these data. We base our criticism on the predominance of corner portfolio allocations in their data, and discuss the correct way to handle these corner allocations keeping in mind the underlying theory being evaluated. Based on these problems alone, their findings are questionable. The second claim we criticize is that "time preferences are not risk preferences." We propose a more careful theoretical statement: if one assumes an additive intertemporal utility function when subjects behave as if they use a non-additive function, then adding risk to time-dated choices will confound inferences about the instantaneous utility function. In fact, non-additivity perfectly explains the pattern of behavior that leads A&S [2012b] to make their striking claim. The third claim we criticize is that, in terms of logistics and ease for the subjects, there is something intrinsically better with the CTB approach than with the BC approach. What is labeled in the CTB approach as one task is in fact a series of tasks varying choice parameters in different ways, generating its own difficulties for inference.

Section 1 provides a review of the methods and findings in A&S [2012a][2012b]. Sections 2, 3 and 4 discuss our criticisms of the three claims, Section 5 briefly reviews open issues to add a broader perspective on this debate about the elicitation of risk and time preferences, and Section 6 concludes.

1. A Brief Review

The basic design of A&S [2012a] is called a Convex Time Budget (CTB) approach. It allows subjects in each choice to choose a portfolio of sooner payoffs and later payoffs. They hold constant the later payoff amount, and vary the sooner payoff amount, but the idea is the same as the BC approach with the addition of allowing a portfolio allocation. For instance, using the "Sample Decision Sheet" in their Figure 1 (p.3340), and ignoring portfolio allocations for the moment, the first choice is between \$20 on a sooner date and \$20 on a later date, the second choice is between \$19 on the same sooner date and \$20 on the same later date, and so on to the last choice between \$14 on the sooner date and \$20 on the later date.[1] For each choice the subject decides on a fraction of their portfolio allocated to the sooner option and the residual fraction allocated to the later option. If the fraction for the sooner option in the first choice is 0.83, as in their example, the subject would receive \$16.60 = \$20 × 0.83 on the sooner date and \$3.40 = \$20 × 0.17 on the later date. If the fraction for the sooner option in the second choice is 0.51, as in their example, the subject would receive \$9.69 = \$19 × 0.51 on the sooner date and \$9.80 = \$20 × 0.49 on the later date.

The innovation here allows for portfolio allocations within the two endpoints of each binary choice. A&S [2012a] find that 70% of the observed choices in this design were at one extreme or the other, where the subject chose to allocate everything to the sooner payoff or the later payoff. The utility function is defined over a single attribute, money, and they find that the parameter that measures the concavity of the utility function is close to one, implying a utility function that is almost linear.

The design of A&S [2012b] is to allow the sooner payoffs to be received with a probability p and the later payoffs to be received with probability P. They consider six {p, P} combinations: {1, 1}, {0.5,

[1] Their subject interface presents these amounts as "exchange rates" between 100 tokens and cash on those dates. So the second choice is actually presented as 100 tokens at an exchange rate of \$0.19 per token or 100 tokens at an exchange rate of \$0.20 per token. It is not clear why the added layer of tokens is needed.

0.5}, {0.5, 0.4}, {0.4, 0.5}, {1, 0.8} and {0.8, 1}. The initial, non-stochastic {1, 1} case is the design in A&S [2012a]. For the other treatments there was an *independent realization* of p and P. These realizations occurred at the same time, so that no preferences for temporal resolution of risk would confound inferences. The key finding is that more subjects allocated their portfolio in the interior when $p<1$ and/or $P<1$. Specifically, in their baseline {0.5, 0.5} treatment only 25% of the choices were at the extremes, compared to 74% for their {1, 1} control.[2] The conclusion drawn from this is that the utility functions over risk and time are not the same. We clarify the possible theoretical interpretations of this conclusion below.

2. Econometrics

Figure 1 shows vividly that the vast majority of the choices in A&S [2012a] involved subjects picking extreme portfolios of sooner and later payoffs. By "vast" we mean 70%, not just a narrow majority, although even a narrow majority would generate the same inferential problems. The choice variable to be modeled, the fraction allocated to sooner payoffs, takes on values between 0 and 100, but with one massive spike at 0 and another clear spike at 100. Since the annual interest rates offered by A&S [2012a] were 0%, roughly 30%, roughly 65%, and then *much* higher, it is no surprise to see many choices to invest with the experimenter.[3] Intermediate values are, as one might expect, clustered around 10% increments, although there are very few values between 51% and 99%. As interest rates offered to the subjects are varied by experimental design, the mix of 0 and 100 choices changes, but the same vast majority of choices are still at the extremes. The top panels of Figures 2 and 3 show the same qualitative

[2] The fraction of choice tasks with a zero interest rate in A&S [2012a] is different from the {1, 1} control of A&S [2012b]. Hence the 70% of corner solutions in A&S [2012a] is close to, but not identical to, the 74% of corner solutions of the {1, 1} control of A&S [2012b].

[3] Figures A1, A2 and A3 in the Appendix show the portfolio allocations for each interest rate offered by A&S [2012a] across treatments with no delay to the sooner option (A1), a front end delay of one week (A2) and a front end delay of five weeks (A3).

pattern for three specific tasks, where we hold constant the annual interest rate offered to subjects.

A. When the Average Is An Illusion

How might one statistically model this type of data? To a trained econometrician, Ordinary Least Squares (OLS) is not the first thing one would think of, since one is inviting inconsistent inferences due to heteroskedasticity.[4] The reason is that the error term in an OLS specification can wander between $\pm\infty$, but the choice data is restricted to [0, 100], and furthermore it is virtually all piled up at 0 or 100. The fact that one uses non-linear least squares (NLLS) does not change this basic point.

The methods proposed by A&S [2012a] focus entirely on explaining the *average* portfolio allocation, even if the average is an illusion that is almost never close to where the data is. They address this issue of the corner allocations in three ways.

The first is to ignore it. The primary NLLS estimates reported in their Table 1, and *all* of their estimates at the individual level, simply assume away the statistical implications of the mass of observed choices at the corners. Figure 4 shows what this model predicts. Figures 2 and 3 show what the prediction is for a specific interest rate offered to subjects in a specific task, and Figure 4 pools these predictions across tasks with different interest rates. As expected, these theoretical predictions are all interior portfolio allocations, which is indeed where the *average* allocation is. But, again, most of the data is actually at the boundaries, resulting in residuals in Figure 5 that are, well, ugly. The theme in the discussion of results is that "it does not matter much" to the results when applying various corrections for these corner choices, so one can safely ignore those messy complications.[5]

---

[4] We are aware that it is the first, and only thing, that many researchers think of. The arguments in Angrist and Pischke [2009; §3.4.2] clearly have to do with estimating *average* effects in randomized control trials, which is not the setting being considered here. There are other concerns with the abuse and defenses of OLS in applied research, but that is beyond our present scope.

[5] A&S [2012a; p.3349] acknowledge that they are forced to ignore these issues for many of the estimates at the individual level, when observed choice data is almost entirely at the corners. In those cases the alternative

The second method is to adapt the criterion function for the NLLS estimator and make "additional distributional assumptions" (p.3342). Specifically, they define a latent variable c^*_t as being equal to the theoretically predicted allocation for some discount rate and utility function and the parameters of that choice (the interest rate and horizon, in particular) *plus* some error term. They then assume (Appendix, page 2) that they "observe $c_t = 0$ if $c^*_t \leq 0$, $c_t = m/(1+r)$ if $c^*_t \geq m/(1+r)$ and $c_t = c^*_t$ otherwise."[6] The first assumption is that the choice will be zero when the latent variable is zero or negative, and the second assumption is that the choice will be 100 if the latent variable indicates it should be 100 or greater. So the first assumption allows for a latent variable, the optimal choice plus a mean-zero error, to be negative. This is theoretically incoherent given their utility specification, which is defined only over money offered in the lab. The second assumption allows for the latent variable to be greater than 100 tokens, which also makes no sense theoretically. We are told that, as "discussed in Wooldridge [2002; ch. 16], c^*_t here does not have an interpretation, but the latent variable vocabulary and associated censoring techniques are applicable to corner solution applications." (A&S [2012a; Appendix, p.3]) But in this instance the latent variable c^*_t has a clear theoretical interpretation, and it was just provided explicitly prior to these censoring assumptions.[7] One might wish that it has no interpretation, to be able to get on with

estimators are simply not defined.

[6] We add brackets to "1+r" in the original since this is obviously intended.

[7] In fact this is not what Wooldridge [2002] says. He notes that censored regression models generally come from two types of categories. "In the first case there is a variable with quantitative meaning, call it y*, and we are interested in the population regression E(y* | x)." (p. 517). The word "quantitative" is crucial here: to use his example, wealth above some level of $200,000 might be top-coded as $200,000. The other case he refers to (p.518) is where y* is some latent variable. In the classical case that inspired the Tobit estimator, we actually only observe expenditures being non-negative: negative expenditures are not even defined conceptually. But in order to allow for the fact that there might be a pile of zero expenditures we assume some latent index y* that *includes an error term* for desired, positive expenditures that allows us to infer the pile at zero as those observations for which the latent index *with* error are non-positive. If the error distribution is known, and it is assumed to be Normal for the Tobit estimator, this is a simple statistical matter. Wooldridge [2002; p. 520] then cautions: "For the [Tobit] statistical model [...] to make sense, the variable y* should have characteristics of a normal random variable." Of course, when y* is very close to zero and the error has modest or large negative values, this latent variable typically takes on negative values: stare again at Figure 1. However, he then goes on to note that for "... corner solution outcomes, we must avoid placing too much emphasis on the latent variable y*. *Most of the time* y* is an artificial construct, and we are not interested in E(y* | x)" (p. 520; emphasis added). This is not a "most of the time

the estimation, but it is critical not to lose touch with the theory at this point. The econometric method is positing something which is theoretically impossible. With mis-specification of that kind, no consistent estimates are possible. Indeed, our Figure 6 shows the estimated portfolio choices using this approach, from the specification in column (1) of A&S [2012a; Appendix Table A1]. Nothing is changed from Figure 4, and it is easy to see why not.

When one inspects what is done with these distributional assumptions, the critical step in their method is to write down an econometric specification where utility curvature (α) and discount rates (δ) *play no role in explaining why one would choose 0% or 100%*, even though these are where the vast majority of the data are! The reason is that the likelihood of the 0% choice is given by the probability that the *error term* leads to a value of c^*_t that is less than or equal to 0, but this probability is affected by the error term and not the part of the criterion function that is affected by α or δ. The error term has two parameters: a mean and a standard deviation. The mean is assumed to be zero, and the standard deviation is typically estimated.[8] But the standard deviation is not affected by α or δ, so these preference parameters are not selected by the least-squares criterion to affect the standard deviation and "explain" the corner probabilities, and hence the predictions of Figure 6 are virtually identical to the original specification in

setting," since the whole point here is to make inferences about structural parameters that define the latent index, the curvature of the utility function and the discount rate. Finally, he warns that, "In corner solution examples, the variable y should be (roughly) continuous when y>0. Thus the Tobit model is not appropriate for ordered responses [...] Similarly, the Tobit model should not be applied to count variables [...]" (p.520). The data in Figure 1, as modeled by A&S [2012a], are ordered variables *and* count variables. Just over 83% of the choices are in 10% increments, and just over 91% of the choices are in 5% increments. These comments on the Tobit model are relevant, as noted below, since A&S [2012a] also use it as a robustness check on their NLLS estimates.

[8] To make matters worse, the error term is *assumed* to have a standard deviation of 1 token, effectively limiting by assumption the extent to which the value of c^*_t is able to go negative. The exposition in A&S [2012a; Appendix] correctly explains the implications of assuming an error with a standard deviation to be estimated, but then they assume it to be 1 without explanation. Their computer code correctly allows explicitly for this parameter, but then just sets it to 1 inside the function evaluation subprogram. The problem is that $\pm 1.96 \times 1$ is never going to be a big number in terms of the dependent variable, tokens defined between 0 and 100. Hence it is no surprise that allowing for corners with vanishingly small probability has no impact on estimates, since the error term has been assumed to be "close" to the average and not allowed, by assumption, to get negative or greater than 100.

Figure 4. Again, this might be a minor matter if the corners were not where the vast majority of the data happen to be piled up. In summary, the problem with the second method is that it does not allow the preference parameters α and δ to affect the probability of the data at precisely the point where most of the data exist.

The third method is to use a two-sided Tobit specification. The idea behind this type of Tobit model is again that there is a latent index which can take on values between $\pm\infty$, but which leads to observed choices that are censored at 0 or 100.[9] Observed spikes at 0 and 100 help explain the determinants of the latent index, where the determinants here are the preference parameters α and β.[10] But this approach only handles the observations literally at 0 or 100, and treats them as binary variables. For an observation at 0, for instance, there is a "probit" specification that determines if the observation is 0 or greater than 0. Compared to the first two methods, this is an improvement since the likelihood of these choices at 0 or 100 corners are at least affected by the preference parameters α and β. But this specification still fails to address the concern that the error term can imply theoretically incoherent values of the latent variable for all interior allocations, again given their utility specification.[11]

[9] The latent index in this case is actually the log of the ratio of consumption sooner and consumption later, so that the index can be expressed as a linear function of variables that allow one to infer back the preference parameters α and β. See equation (6) of A&S [2012a]. This log transformation is used so that standard computer packages for censored regression can be utilized.

[10] A further improvement would be a "hurdle" specification in which the binary process determining a corner can be separate from the process determining interior choices. This specification is common in health economics, where it is used to capture the idea that the factors that cause someone to seek medical care from a doctor are distinct from the factors that cause the doctor and patient to decide how much to spend. In this case going to the doctor is the hurdle that must be passed before expenditures would be positive (e.g., Coller et al. [2002]). It has also been used in experimental economics to examine free-riding behavior in a public goods setting: in this case the subject has to decide whether to contribute any amount at all, and only then does the process determining the positive contribution level apply (e.g., Botelho et al. [2009]). In many applications the insights from hurdle models are significantly different from those from Tobit models.

[11] See the discussion in Wooldridge [2002], cited above, about the correct interpretation of the Tobit estimator. One econometric solution for these problems, if one insists on trying to estimate the preference parameters with canned statistical packages, is what is known as a "truncated regression." In this case the error term is truncated and re-normalized at the upper and lower bounds, so that the latent variable cannot take on incoherent values. The problem with this approach is that canned routines assume a normal error, which gets re-normalized "into" the interior. This can lead to inconsistent results if the data-generating process causing corners

B. Modeling the Distribution, Not Just the Average

As it happens, there are several simple ways to estimate portfolio choice models of this kind, where the statistical objective is to make inferences about structural preference parameters and the whole distribution of observed data, and not the average of two corner modes. The most natural approach is to assume that subjects choose one of the 101 percentage point allocations between 0 and 100, based on the present value of the utility of the alternative. Let the present value of the alternative be evaluated with an exponential discounting model, so that the discount factor is $D_t^E = 1/(1+\delta)^t$ for horizon t, and the discount rate is then $d_t^E = \delta$. Or one could use a quasi-hyperbolic discounting model, so that the discount factor is $D_t^{QH} = 1$ if $t = 0$ and $D_t^{QH} = \beta/(1+\delta)^t$ if $t > 0$, where $\beta<1$ implies quasi-hyperbolic discounting and $\beta=1$ is exponential discounting. The discount rate is then $d_t^{QH} = [\ \beta/(1+\delta)^t\ ]^{(-1/t)} - 1$ for $t > 0$. If we denote $\tau \in [0,1]$ to be the allocation $\{\tau c_t, (1-\tau)c_{t+k}\}$, then we evaluate the (additive) intertemporal utility function $U_\tau = U(c_t, c_{t+k}, \tau) = D_t\, u(\tau c_t) + D_{t+k}\, u((1-\tau)c_{t+k})$. We can then define the latent index

$$\nabla U = \exp(U_\tau)/\ [\ \exp(U_{0\%}) + \exp(U_{1\%}) + \ldots + \exp(U_{100\%})\ ]$$

for the specific choice τ that was observed. This latent index is the familiar multinomial logit probability of observing the choice τ.[12]

This specification has no "corner" solution issue demanding special attention, providing there is some variability in corners over the sample. If there are more observations at τ=0% and τ=100% than other values of τ, then that raises no issues at all. The maximum likelihood task is to find the parameters defining the discounting function and utility function so as to maximize the probability of the observed

---

is different from the data-generating process causing interior choices, as it is here (e.g., weakly convex utility can lead to corners, whereas strictly concave utility implies interior choices, *ceteris paribus* the discount rate). Davidson and MacKinnon [1993; p. 534ff.] provide excellent expositions of these issues.

[12] It is not appropriate to use an *ordered* logit specification in this specification, since the ordering of the latent index ∇U with respect to the ordering of the choice τ is not generally increasing or decreasing $\forall\tau$. This is different from the A&S [2012a] specifications, where the dependent variable is the number of tokens allocated to the sooner option, which is cardinal and ordered.

choices. We estimate this model using the A&S [2012a] data and for direct comparability to the various methods used by A&S [2012a], we assume the same (homegrown) endowment of $7.05 that they estimated from the survey question they asked their subjects.[13]

If we assume a linear utility function this specification implies an average discount rate estimate of 28.4%, with a 95% confidence interval of 15.4% ↔ 41.4%. If we allow for a non-linear utility function this specification actually implies a *convex* utility function, with a RRA estimate of -0.36 with a 95% confidence interval of -0.45 ↔ -0.26, and a higher discount rate estimate of 60.1% with 95% confidence interval 37.6% ↔ 82.7%. Why do we find convex utility, which is almost never observed in comparable lab settings? The reason is that the model has to account for why the representative agent wants to pick the 0% corner solution in the majority of cases, but not forgetting that there is a second important mode at 100%. In a maximum likelihood sense, these are the two choices out of the 101 possible allocations that are the most important to explain. Of course, the model seeks to explain all choices with the highest joint probability possible. Figure 7 shows the predictions of this specification when we constrain to a linear utility function, and when we allow a convex utility function. It is apparent that the linear utility function handles the spike at 0% easily enough, but assigns too much mass to interior solutions, and far too little to the spike at 100%. The convex utility function is bunched right at the boundary allocations, just as the data is, with no mass to interior solutions that is visible here. It assigns a bit too much mass to very small and very large allocations, but generally captures the right pattern of the distribution of observed choices.[14]

---

[13] There are some problems identifying these endowments econometrically, whether or not they are assumed to be the same in the sooner and later periods. The reason is that different endowments affect inferred utility curvature, which in turn affects inferred discount rates. This problem is particularly serious when one tries to estimate different endowments for sooner and later periods, since this allows different utility curvature in sooner and later periods solely because of the different endowments. The fact that numerical procedures converge on some local optima says nothing about identification *per se*. There are deeper issues involved here, to do with the extent of asset integration between lab prizes and extra-lab wealth or income; we discuss these later in section 5.

[14] We applied the same model to individual choices. For 7 of the 97 subjects there was no variation at all in the observed (corner) choice, and for 21 subjects there was only 1 corner choice at 0% and 44 corner choices at 100%. For these subjects there is simply not enough variation in choices to estimate any model. For the 53 subjects for which we could estimate the model, the distribution of estimates is broadly consistent with the

To understand the economics behind this outcome, let us consider the first five tasks presented by A&S [2012a; Table 1, p.3338] under various assumptions of utility curvatures. In these decisions the subject chooses a portfolio between money today and money in 35 days time. Our Table 1 shows the predicted optimal allocation of income, using the estimated risk attitudes and discount rates from AHLR [2008a] and then the comparable estimates from A&S [2012a; Table 2, column 1, p. 3346]. The former study found an annual discount rate δ of 10% with a CRRA coefficient α of 0.74, using the specification $u(x) = x^{1-\alpha}/(1-\alpha)$, and the latter study found a discount rate of 30% with a CRRA of 0.08. The differences in the studies are clear, but not important for present purposes; just view the first study as one finding relatively concave utility and low discount rates, and the latter study as one finding very little concavity of utility and higher discount rates. The key point of comparison is the prediction, conditional on the two sets of values of δ and α, of the optimal fraction to be allocated to the sooner amount.[15] The concave utility subjects in Panel A of our Table 1 should choose an interior solution in every case, and a "deeply interior" solution close to 50%. The mildly concave utility subjects in Panel B of our Table 1 should choose an interior solution as well in every case, but one that declines as we move down the choices. The observed choice pattern from A&S [2012a] clearly shows a decline if one just looks at the average contributions, so it would seem apparent that the second set of preference parameters does a better job of explaining the data compared to the first.

But the predictions in Panel B are that *each and every* subject, armed with these "representative agent" preferences, would choose interior allocations *every* time. If the observed 42% *average* in choice set 1 actually consisted of 42 subjects picking exactly 0% and 58 subjects picking exactly 100%, the average would look like 42% but would wildly distort the theoretical prediction. In fact, the observed data are

---

statements about the representative agent: convex utility, and annual discount rates between the offered interest rates of 66% and 113%.

[15] For simplicity we assume zero endowments and exponential discounting.

close to that extreme, as Figure 2 shows for one illustrative example. One certainly would allow the theory some minor deviations either side of the 43% prediction in the bottom panel of Figure 2, but one can hardly compare the two panels of Figure 2 and claim that the *qualitative* behavior predicted is consistent with the observed data.

In Panel B of Table 1 we also show the parameters and predictions from A&S [2012a] for two other choice tasks. These were the two additional tasks with the sooner amount of money available in the present and for which the smallest interest rate was provided. For choice task 1 the interest rate was 65%, so any subject that selected 0% of the sooner allocation with a linear utility function would have revealed a discount rate between 0% and 65%. The respective upper bounds for choice tasks 6 and 11 are 30% and 21%. The observed data for these three tasks, displayed in the top panels of Figures 2 and 3, allow us to see below what is going on with the econometric methods of A&S [2012a]. Three things are again apparent from Figure 3, as with Figure 2: the vast majority of the observed data is at the corners, the parameters α=0.08 and δ=0.3 generate striking interior predictions and *no* predictions of corner choices whatsoever, and these predictions match the *average* of the observed data quite closely (see our Table 1). Thus, the estimated model explains average responses quite well, but not the underlying distribution of responses.

Panel C of Table 1 shows the effect of assuming effectively linear utility for the representative agent with the same discount rate estimated by A&S [2012a]. The optimal choice is always 0%, as one would expect given the assumed discount rate of 30% and the implied interest rates. Of course, 0% is one of the corner solutions. If these predictions are compared to the *average* choice, they look terrible. But when we reflect on Figure 2, and the fact that the *modal* choice is in fact at the 0% corner, it is not at all obvious that this is a poor set of estimates. The convex utility model would make the same prediction as the linear utility model with an assumed discount rate of 30% and implied interest rate of 65%.

Only a convex utility function can explain the choice pattern in Figure 7 under the assumption of homogeneous representative agents. Concave utility cannot help here, since it forces (virtually) all optimal choices into the interior, and linear utility allocates too few choices to later payoffs. The convex utility function, however, encourages choices at either extreme, *ceteris paribus* the discount rate δ, because marginal utility is increasing over income. It will also lead to higher discount rates because of Jensen's inequality and move optimal choices towards later payoffs. In other words, a convex utility specification with a positive discount rate predicts some choices at either extreme, which is what we generally observe. By focusing on estimating a structural model of averages that have virtually zero mass in the data (near the average, not just at the average), one misses the simple economics of behavior here.

Of course, does anyone believe that subjects in laboratory experiments such as these exhibit convex utility functions? Even the claim of linear utility is tenuous, given the way that the utility function has been modeled.[16] The vast bulk of the evidence, for these stakes and this way of modeling utility functions, is for concave utility (e.g., see Harrison and Rutström [2008] for a survey). Agreed, this evidence comes from experiments examining risk preferences, and A&S [2012b] claim that these have different utility functions than utility over non-stochastic outcomes. We question that claim on theoretical grounds below, and Cheung [2012] is unable to replicate it in a more familiar BC experimental design.

All of the econometric problems with A&S [2012a] are inherited by A&S [2012b], quite apart from any other confounds from theory that the latter design has.

[16] A&S [2012a; p.3334] casually claim that recent calibration findings support their conclusion of approximately linear utility: "Rabin [2000] shows that under expected utility theory, individuals should have approximately linear preferences for small stakes outcomes, such as those normally used in time preference experiments." But Rabin [2000] shows these things when one assumes perfect asset integration of those small stakes with a wide range of extra-lab wealth levels. This is not the structural model A&S [2012a][2012b] estimate, so their claim of support for their findings confuses the one-line punchline rhetoric of the calibration literature with the logic behind it.

C. Time and Risk

Figure 8 displays the raw portfolio choice pattern from the experiments of A&S [2012b]. The control here, shown in the top left panel with {p,P}={1,1}, is the same design used in A&S [2012a], with two qualifications. The first qualification is that only choices with a sooner payment in 7 days are used in A&S [2012b], and with horizons of 28 and 56 days. The second qualification is that there are many more choices with a zero interest rate in A&S [2012b], in which the same amount of money was offered in the sooner and later period: 960 out of 6,720 choices, or 14.3% in A&S [2012b], compared to only 97 out of 4,365 choices, or 2.2% in A&S [2012a].[17]

The other panels in Figure 8 show the effects of allowing for risk. As we move from {p, P} of {1, 1} to {1, 0.8} in the first two top panels, we see a striking shift towards subjects choosing the sooner tokens. The reason is simple, if the subject is modestly risk averse: the later payoff is now risky, and the sooner payoff remains certain. Moving along to the panel with {p, P} = {0.8, 1}, in the top right corner, we see the opposite shift: now the tendency under certainty to select the later option is exacerbated, since the sooner option is risky and the later option is certain. Again, a simple explanation if the subject is risk averse.

The bottom panels in Figure 8 are where the real action occurs in A&S [2012b]. Here we see a general tendency, compared to the certain case in the top, left corner, for more interior choices. In the choices with *zero* interest rate and {p, P} = {0.5, 0.5}, where the sole motivation for diversification can be correlation aversion, discussed below, we indeed see over 60% of the choices exactly at 50%, and the next most important mode at 100% sooner allocations.[18] Understanding this behavior requires a careful statement of the appropriate theory.

---

[17] Appendix Figures A4 and A5 show the choices with positive interest rates and zero interest rates, respectively, since the latter distort interpretations in some cases.

[18] Figure A5 shows these results, in the bottom left corner.

3. Risk Preferences Are Not Time Preferences

When one assumes an additive intertemporal utility function, the curvature of the instantaneous utility function determines the manner in which risky prospects over time-dated outcomes are traded off. Indeed, as is well known, it "overdetermines" this tradeoff, since relative risk aversion has to be the same as the inverse of the intertemporal elasticity of substitution, generating calibration headaches for many decades. But if one allows a non-additive intertemporal utility function, it is also well known that one can tease apart a-temporal risk preferences from temporal risk preferences. That is, one can then separately identify the curvature of the level functions of utility and the curvature of the expansion path of utility. The implication is that attitudes towards risk over differently-time-dated amounts can be very different from attitudes towards risk over same-time-dated amounts. This is known in the literature as "correlation aversion," following Keeney [1973], Richard [1975] and Epstein and Tanny [1980], who gave it this name.[19]

The idea is quite simple. Assume for the moment that the discount rate is zero and imagine payoffs X>x payable sooner in time, and payoffs Y>y payable later in time. Define a lottery α that is a 50% chance of {X, y} and a 50% chance of {x, Y}, spreading one's eggs over time. Then define a lottery β that is a 50% chance of {X, Y} and a 50% chance of {x, y}, the feast or famine options. If someone is indifferent between α and β for all x, X, y, Y that meet the ordering constraints noted, then we say that they are correlation neutral. If someone prefers α to β we say that they are correlation averse, and if they prefer β to α we say that they are correlation loving. The concept is akin to the idea of a-temporal risk aversion deriving from instantaneous utility curvature, where one prefers to reduce a-temporal variability in outcomes. The literature establishes the remarkable result that the only (separable) intertemporal utility

[19] Non-additivity also arises from "habit formation" models, such as those considered by Constantinides [1990] to address the equity premium puzzle and by Davidoff, Brown and Diamond [2005] to address the annuitization puzzle.

function that characterizes a correlation neutral agent is an additive function, and that this is "if and only if." So if one casually assumes additivity, because everyone else does, then one is committing to assuming correlation neutrality.

The implication for the claim that "risk preferences are not time preferences" is immediate. If the intertemporal utility function that subjects use is actually non-additive, then risk preferences over time streams of money need to be sharply distinguished from risk preferences over a-temporal payoffs. In effect, there are two possible types of risk aversion when one considers risky choices over time, not one. To be more precise, if one gives subjects choices over differently-time-dated payoffs, which is what A&S [2012b] did, one sets up exactly the thought experiment that *defines* correlation aversion. They compare behavior when subjects make choices over time-dated payoffs that are not stochastic with choices over time-dated payoffs that are stochastic, and observe different behavior. In the former case virtually all choices in their portfolios were at extreme allocations, either all payoffs sooner or all payoffs later; in the latter case they observed more choices in which subjects picked an interior mix of sooner and later payoffs, diversifying intertemporally. Evidence that subjects behave differently, when there is an opportunity for correlation aversion to affect their choices compared to a setting in which it has no role, is evidence of correlation aversion. It is not necessarily evidence for the claim that there is a "different utility function" at work when considering stochastic and non-stochastic choices. We do not rule the latter hypothesis out, but there is a simpler explanation well within received theory.[20]

Evidence for correlation aversion in experiments is provided by AHLR [2011c], who also provide extensive cites to the older literature. Cheung [2012; Appendix A.2] proves nicely that correlation aversion

---

[20] We assume in AHLR [2008a] that individuals have stable utility functions between period t and period t+k, and furthermore that their subjective beliefs about those utility functions are stable. We have offered considerable data for the former assumption of stability in Harrison, Johnson, McInnes and Rutström [2005] and AHLR [2008b]. Although we are skeptical of the evidence in Loewenstein, O'Donoghue and Rabin [2003], we accept that the ability to "project" utility functions into the future is an open and important question.

provides an immediate explanation for the observed behavior in A&S [2010b].[21] He refers to it as a motive for intertemporal diversification, which is a valid way of characterizing the idea of correlation aversion. Just as a-temporal risk aversion encourages mean-preserving reductions in the variability of a-temporal payoffs (imagine lotteries defined solely over x and X or defined solely over y and Y), correlation aversion or intertemporal risk aversion encourages mean-preserving reductions in the variability of the time stream of payoffs (imagine lotteries α and β defined above over x, X, y and Y).

Hence, when A&S [2012b] claim that "risk preferences are not time preferences," one can restate this more carefully as "a-temporal risk aversion is not the same as intertemporal risk aversion," and of course that is correct whenever there is a non-additive intertemporal utility function.

One further comment is justified based on the reliance on EUT when comparing utility curvature across tasks with and without risk, the basis of the claims of A&S [2012b]. In any task that involves risk it is conceivable that part of the risk attitude is captured through transformations of the probabilities, as in Rank Dependent Utility (RDU) or Cumulative Prospect Theory (CPT), and not only through utility curvature. If RDU or CPT is the true model, then estimating preferences relying on EUT can lead to over- or underestimations of utility curvatures.

To illustrate with a specific example, consider the first-order condition for utility maximization when there are probabilities p and P of payment in the sooner and later periods, respectively. If one assumes an additive intertemporal utility function then it can be shown that $D_t \; p \; u'(c_t) \; / \; D_{t+k} \; P \; u'(c_{t+k}) = (1+r)$. If p=P the probabilities cancel, so one should see the same choices, and infer the same discount factor and discount rates, whatever the specific value of the two probabilities. In particular, the cases p=P=1 and p=P=0.5 are two studied by A&S [2012b] and Cheung [2012].

Let there be a power utility function $u(x) = x^\sigma$, and suppose that one estimates $\sigma = 0.5$ from the

[21] We first raised this point with A&S at a conference that we organized in Denmark in June 2010, and explained it to Cheung in 2011.

observed data when assuming EUT and the choice options are stochastic. Then we have a noticeably concave utility function. The utility values of the $16.60 and $3.40 payoffs in the earlier example on page 3 are 4.07 and 1.84, and their ratio is 2.21. If the individual followed RDU in making decisions under risk and had a pessimistic probability weighting function, we would have a true σ closer to 1 than 0.5, since the pessimistic probability weighting function would explain part of the aversion to take the risky option. Say that we then have a true σ = 0.7; the ratio of the marginal utilities of the same payoffs would instead be 3.03 = 7.15/2.36. These two cases, σ = 0.5 and σ = 0.7, would lead to very different inferences about the discount factor and discount rates. But using only one non-unitary probability is not sufficient to properly identify general probability weighting functions; at least three points on the cumulative density function would need to be assessed, without overly strong parametric assumptions. With such a restricted design it is necessary to temper the conclusions about relative curvature of the utility function under certainty and uncertainty, cautioning that the conclusions holds only if EUT is the true model.[22]

4. Convex Budget Approaches Are Better Than Binary Choice Approaches

There are two aspects to the comparison of CTB and BC. The first is theoretical and the other is logistical and behavioral. The theoretical foundation for the comparison was introduced by Cubitt and Read [2007], who pointed out that with a linear a-temporal utility function the individual would rationally pick corner solutions, but that it is in principle impossible to identify discount rates if the agent has a non-linear a-temporal utility function. The reason is that a subject might then want to select an interior portfolio allocation between sooner and later payoffs, but is constrained by the binary choice procedure.

---

[22] One might argue that allowing for a probability weighting function ω(·) would just mean that (4) is replaced by $\{ D_t\, \omega(p)\, u'(c_t) \,/\, D_{t+k}\, \omega(P)\, u'(c_{t+k}) \} = (1+r)$. Hence if p = P then ω(p) = ω(P), and these probabilities again drop out. This is correct, but neglects the fact that one infers a different u(·) when there is probability weighting compared to when there is none, and it is the difference in the marginal utilities that matters for inference about time preferences.

While theoretically valid, the empirical magnitude of the problem has not been evaluated to invalidate the BC. The problem for inferences about discount rates arises only if intertemporal utility functions are sufficiently non-linear (and concave) and the interest rate intervals in the BC are sufficiently large, and we do not yet know how concave it has to be in order to be problematic. One simple solution to this problem is to reduce the interest rate intervals in the BC and increase the number of decision tasks. Another solution is to use methods that elicit point estimates, but that approach introduces other problems such as payoff dominance (Harrison [1992]).

The second basis for comparison is the logistical ease of delivering the task, and the comprehension and motivation among participants in responding truthfully. In A&S [2012a] the CTB is presented as a "single, simple instrument." While it is true that the approach uses the same type of task to identify both discount rates and utility curvatures, it can hardly be considered a single instrument. In A&S [2012a] the task involves a series of 45 choice tasks that vary the ratio of payoffs, the time dating of the sooner payment, and the time dating of the later payment. Since task instructions only have to be given once this may indeed be a reduction in cognitive costs to subjects compared to when two task instructions are given, as in AHLR [2008a]. However, it is also possible that not breaking the monotony of making a series of decisions with something new can lead to boredom.

It is also not clear that the task is simpler because it is just one task.[23] In the A&S [2012a; Figure 1, p. 3340] example the subject chooses allocations between $16.60 sooner or $3.40 later, then between $9.69 sooner and $9.80 later, then between $7.74 sooner and $11.40 later, then between $3.36 and $15.80 later, and finally between $1.96 sooner and $17.20 later. The earlier BC approach of Coller and Williams [1999] and Harrison, Lau and Williams [2002] always held one of these sooner or later dollar amounts constant,

[23] There are also some semantic games here. In AHLR [2011c] we presented one task that presented subjects with binary choices over time-dated lotteries. By selecting special cases of the parameters of this task we could have generated each of the other two tasks in that overall design and claimed to have one task.

to make it cognitively easier for subjects to see what the tradeoffs were on the temporal dimension. This CTB design varies two things at once for a given horizon, the size of the sooner and later prizes *and* their ratio, and that could be confusing.

We completely agree that what is complex to one subject might be simple to another, and *vice versa*, depending on the nature of the task and the background of the subject. But we prefer not to build in complexity if it is not needed, recognizing that the modest price for that is the need to have two choice tasks instead of one. Fancier mousetraps can be useful, but sometimes it is more efficient and reliable just to stick to the old mousetrap and buy two of them.

5. Open Issues

Reliably inferring risk and time preferences is not easy. Despite the progress of recent years, we believe that there are several important open issues.

First, we need to allow for alternative models of decision-making over risk for some decision-makers in some settings. The identification of non-standard models of risk preferences demands careful attention to the tasks given to subjects, and is not something that we believe can be safely "folded in" with some other task, as proposed by A&S [2012a][2012b].[24]

Second, the implications for inferred risk attitudes of worrying about asset integration are more subtle than recent controversies over calibration might suggest. Proper identification of the extent to which subjects in experiments integrate the prizes in tasks with "outside wealth" demands unique data,

[24] A subtle complication, making this even harder to draw reliable inferences, is the specific modeling of the "certainty effect" or "boundary effects" in the literature, which is exactly the well-known theoretical point A&S [2012b] make in the context of decisions over time-dated money. Schmidt [1998] reviews the larger literature, and proposes an axiomatization of the certainty effect that relaxes the familiar independence axiom of EUT. So claims about "risk preferences not being time preferences" can be interpreted as saying that there is a certainty effect which also happens to apply to intertemporal allocations, and which can be axiomatized as a violation of the independence axiom.

and does not always lead to the conclusion that subjects have to be risk neutral over small stakes (e.g., Andersen, Cox, Harrison, Lau, Rutström and Sadiraj [2011]).

Third, the discount rates that are implied over monetary prizes and over consumption flows can be quite different, as stressed by Cubitt and Read [2007]. We do not take the view that the only possible or interesting argument of a utility function is a consumption flow, but we are certainly interested in such flows. There are methods for allowing explicitly for the relationship between monetary prizes and consumption flows, as we illustrate at length in AHLR [2008a]. There is a need for comparable experiments examining risk and time preferences over real flows, although "sips of juice" and equally contrived examples of real effort are not what we find convincing.

Fourth, our approach to identification of discount rates defined over utility has always made one assumption we find problematic: that the a-temporal utility function the subject exhibits today is the same a-temporal utility function the same subject applies to evaluate future monetary prizes or consumption flows. In behavioral terms, we assume away any "projection bias," as noted earlier, and should instead use the subjectively expected utility for the future self. We are not aware of an instrument to reliably identify the latter concept, as yet.

Finally, we have to move away from assuming additive intertemporal utility functions, for reasons clear enough from our earlier comments on the obvious confound it implies for the conclusions of A&S [2012b]. The good news here is that instruments are developed to control for this confound.

6. Conclusions

To infer time preferences reliably one has to get the theory, experimental design, and econometrics straight.

Figure 1: Observed Portfolio Choices from the Andreoni and Sprenger [2012a] Experiments

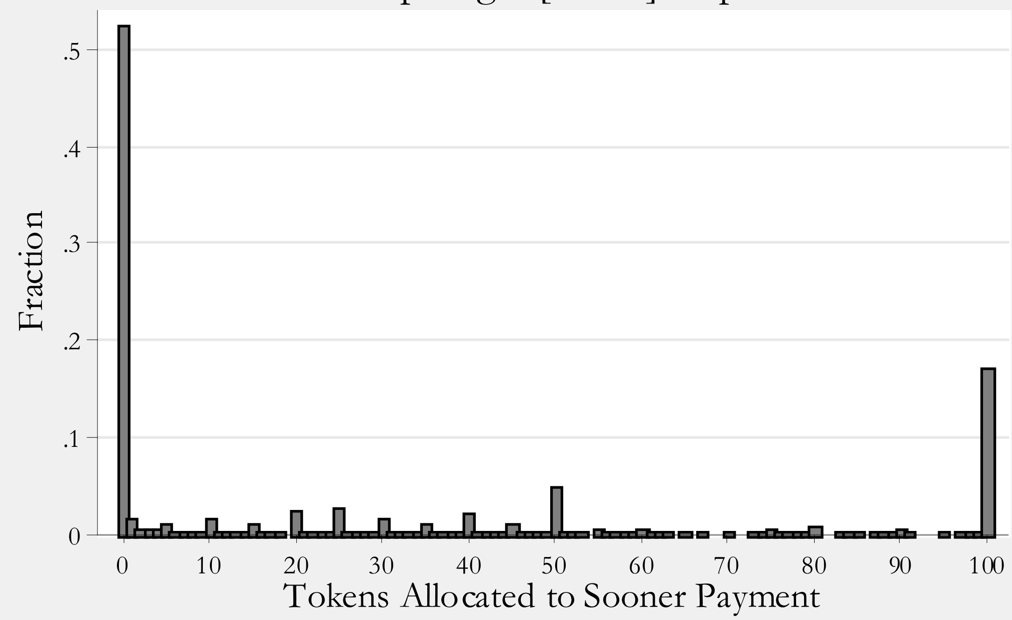

Figure 2: Observed and Predicted Portfolio Choices from the First Choice Task

Later budget offers an annual interest rate of 65%

Prediction with α = 0.08 and δ = 0.3

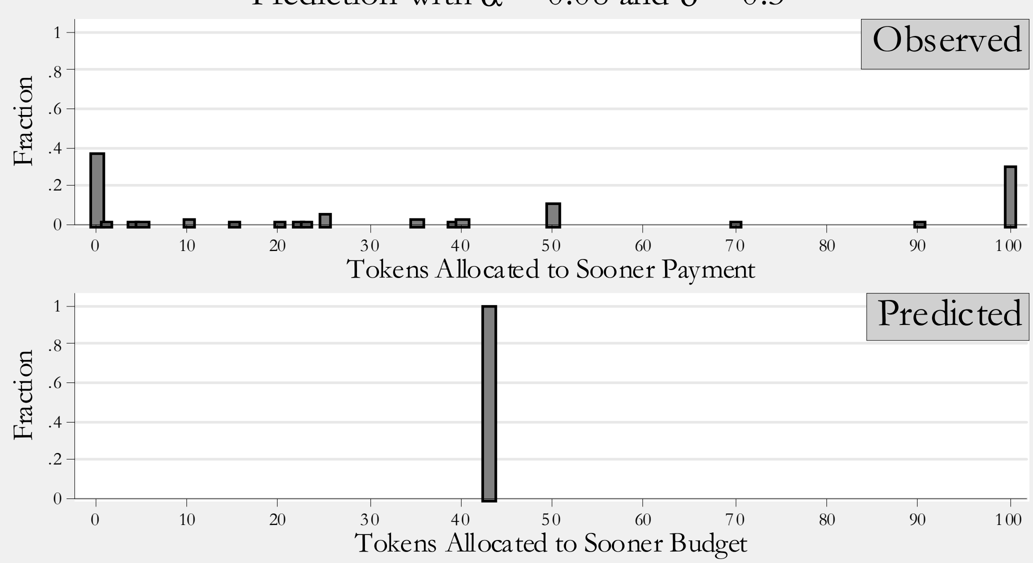

Figure 3: Observed and Predicted Portfolio Choices
from the Sixth and Eleventh Choice Tasks

Later budgets offers annual interest rate of 30% and 20%, respectively
Prediction with α = 0.08 and δ = 0.3

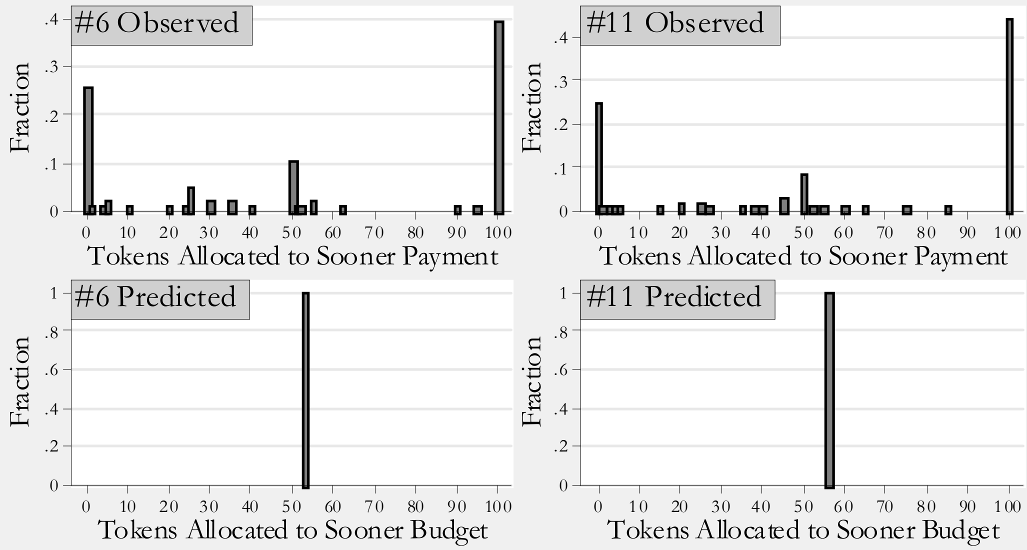

Figure 4: Estimated Portfolio Choices from the
Andreoni and Sprenger [2012a] Econometrics

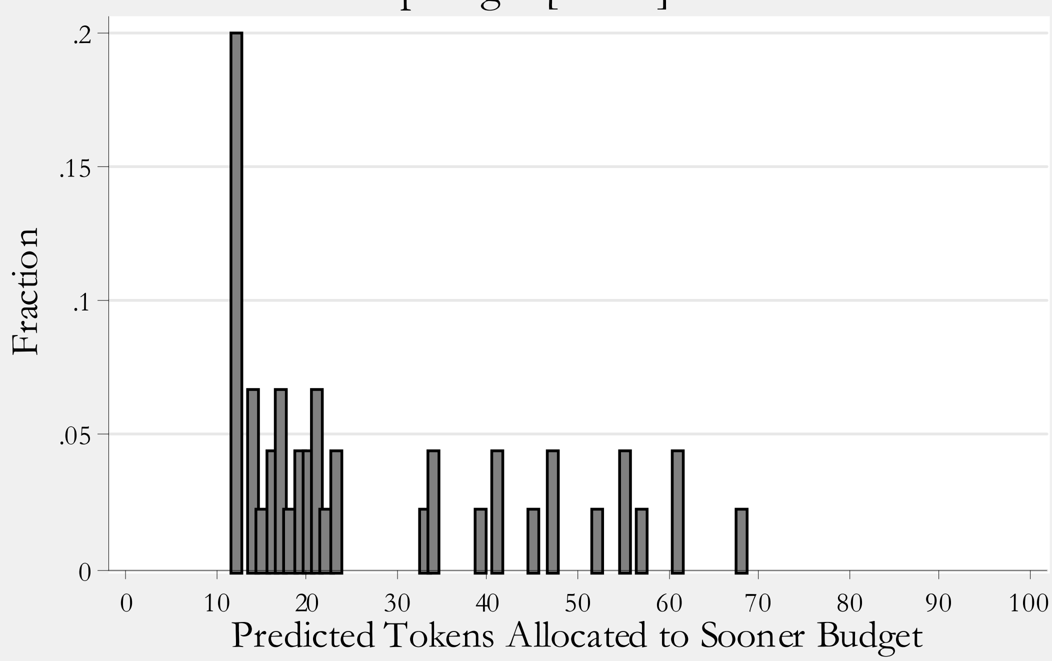

Figure 5: Residuals of NLS Estimation

Shapori-Francia Test for Null Hypothesis of Normality: *p*-value = 0.00001

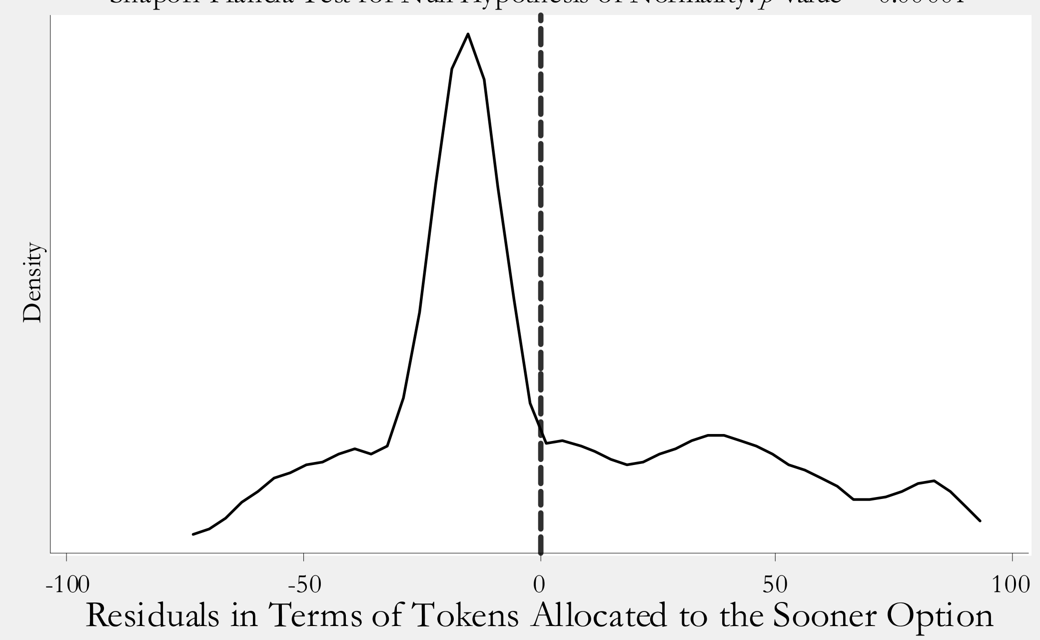

Figure 6: Estimated Portfolio Choices When
Corner Allocations are Modelled In the Error Term

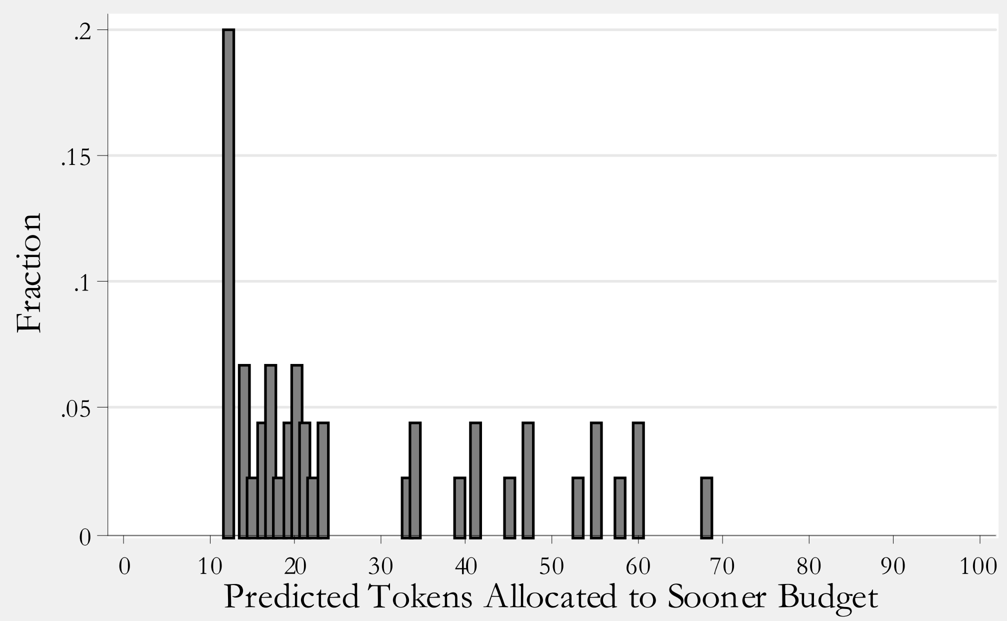

Figure 7: Observed and Predicted Choices from Representative Agent Model of the Distribution

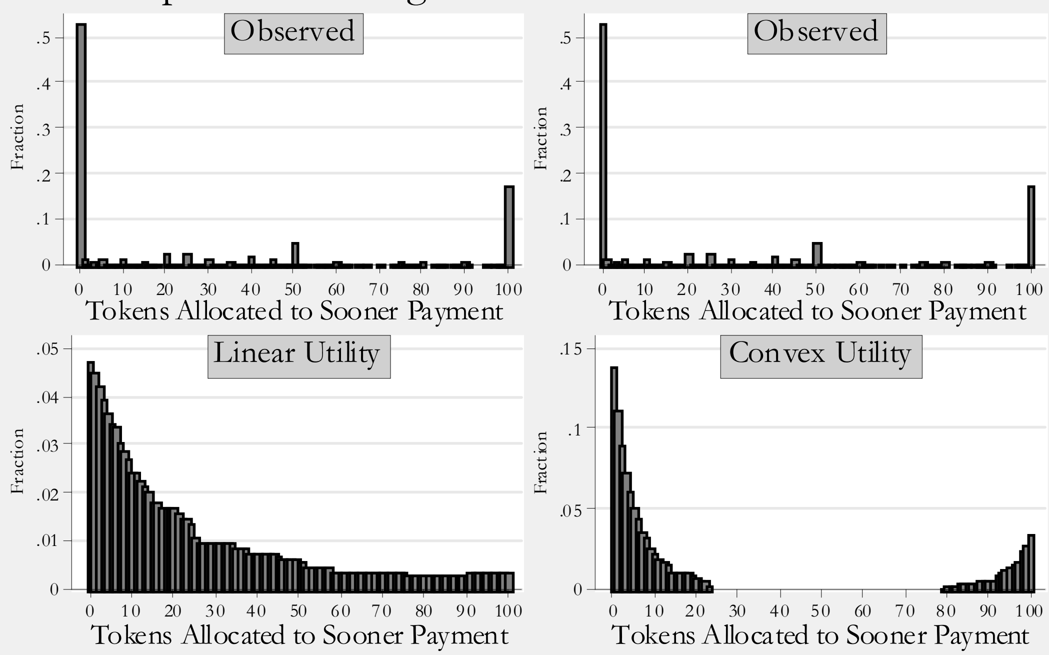

Figure 8: Observed Portfolio Choices from the Andreoni and Sprenger [2012b] Experiments

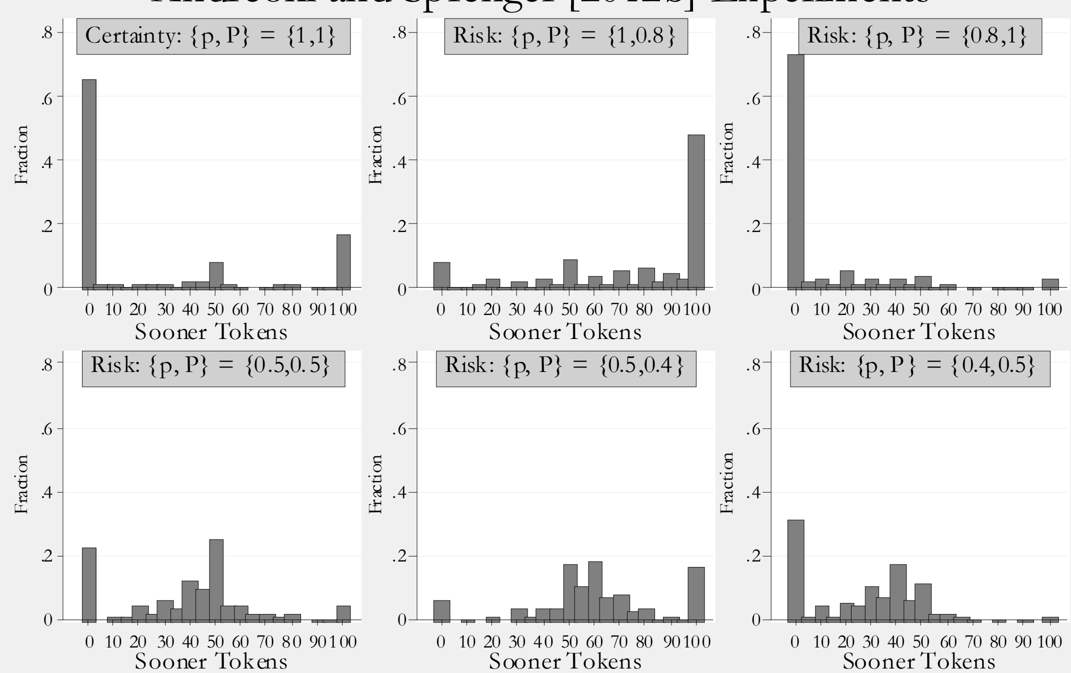

Table 1: Illustrative Choice Predictions

Parameters from the first 5 choice tasks of Andreoni and Sprenger [2012a; Table 1]
Annual discount rate δ. CRRA coefficient α, using $u(x) = x^{1-\alpha}/(1-\alpha)$.

δ	α	Choice Set	Option A Now	Option B Later	Horizon in Days	Annual Interest	Optimal Share	Observed Share
A. Estimates of δ and α from Andersen, Harrison, Lau and Rutström [2008a]								
10%	0.74	1	$19	$20	35	65%	50%	42%
10%	0.74	2	$18	$20	35	164%	49%	31%
10%	0.74	3	$16	$20	35	529%	48%	16%
10%	0.74	4	$14	$20	35	1301%	47%	10%
10%	0.74	5	$20	$25	35	529%	48%	11%
B. Estimates of δ and α from Andreoni and Sprenger [2012a]								
30%	0.08	1	$19	$20	35	65%	43%	42%
30%	0.08	2	$18	$20	35	164%	29%	31%
30%	0.08	3	$16	$20	35	529%	10%	16%
30%	0.08	4	$14	$20	35	1301%	2%	10%
30%	0.08	5	$20	$25	35	529%	10%	11%
30%	0.08	6	$19	$20	70	30%	51%	53%
30%	0.08	11	$19	$20	98	21%	58%	56%
C. Estimates Reflecting 30% Discount Rate and Linear Utility								
30%	0.001	1	$19	$20	35	65%	0%	42%
30%	0.001	2	$18	$20	35	164%	0%	31%
30%	0.001	3	$16	$20	35	529%	0%	16%
30%	0.001	4	$14	$20	35	1301%	0%	10%
30%	0.001	5	$20	$25	35	529%	0%	11%

References

Andersen, Steffen; Cox, James C.; Harrison, Glenn W.; Lau, Morten; Rutström, E. Elisabet, and Sadiraj, Vjollca, "Asset Integration and Attitudes to Risk: Theory and Evidence," *Working Paper 2011-17*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, 2011.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet, "Eliciting Risk and Time Preferences," *Econometrica*, 76(3), 2008a, 583-619.

Andersen, Steffen, Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet, "Lost in State Space: Are Preferences Stable?" *International Economic Review*, 49(3), 2008b, 1091-1112.

Andersen, Steffen; Harrison, Glenn W., Lau, Morten I., and Rutström, E. Elisabet, "Discounting Behavior and the Magnitude Effect," *Working Paper 2011-01*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, 2011a.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet, "Discounting Behavior: A Reconsideration," *Working Paper 2011-03*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, 2011b.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten, and Rutström, E. Elisabet, "Multiattribute Utility Theory, Intertemporal Utility, and Correlation Aversion," *Working Paper 2011-04*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, 2011c.

Andreoni, James, and Sprenger, Charles, "Estimating Time Preferences from Convex Budgets," *American Economic Review*, 102(7), December 2012a, 3333-3356.

Andreoni, James, and Sprenger, Charles, "Risk Preferences Are Not Time Preferences,"*American Economic Review*, 102(7), December 2012b, 3357-3376.

Angrist, Joshua D., and Pischke, Jörn-Steffen, *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton: Princeton University Press, 2009).

Botelho, Anabela; Harrison, Glenn W.; Pinto, Lígia M. Costa, and Rutström, Elisabet E., "Testing Static Game Theory with Dynamic Experiments: A Case Study of Public Goods," *Games and Economic Behavior*, 67, 2009, 253-265.

Cheung, Stephen L., "Risk Preferences Are Not Time Preferences: Comment," *IZA Discussion Paper 6762*, Institute for the Study of Labor (IZA), Bonn, July 2012.

Coller, Maribeth; Harrison, Glenn W., and McInnes, Melayne Morgan, "Evaluating the Tobacco Settlement: Are the Damages Awards Too Much or Not Enough?"*American Journal of Public Health*, 92(6), June 2002, 984-989.

Coller, Maribeth, and Williams, Melonie B, "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 1999, 107-127.

Constantinides, George, "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98(3), 1990, 519-543.

Cubitt, Robin P., and Read, Daniel, "Can Intertemporal Choice Experiments Elicit Time Preferences for Consumption?" *Experimental Economics*, 10(4), 2007, 369-389.

Davidoff, Thomas; Brown, Jeffrey R., and Diamond, Peter A., "Annuities and Individual Welfare," *American Economic Review*, 95(5), December 2005, 1573-1590.

Davidson, Russell, and MacKinnon, James G., *Estimation and Inference in Econometrics* (New York: Oxford University Press, 1993).

Epstein, Larry G., and Tanny, Stephen M., "Increasing Generalized Correlation: a Definition and Some Economic Consequences," *Canadian Journal of Economics*, 13, 1980, 16-34.

Harrison, Glenn W., "Theory and Misbehavior of First-Price Auctions: Reply," *American Economic Review*, 82(5), December 1992, 1426-1443.

Harrison, Glenn W.; McInnes, Melayne; Johson, Eric, and Rutström, Elisabet, "Temporal Stability of Estimates of Risk Aversion," *Applied Financial Economics Letters*, 1, 2005, 31-35.

Harrison, Glenn W.; Lau, Morten I., and Williams, Melonie B., "Estimating Individual Discount Rates for Denmark: A Field Experiment," *American Economic Review*, 92(5), December 2002, 1606-1617.

Harrison, Glenn W., and Rutström, E. Elisabet, "Risk Aversion in the Laboratory," in J.C. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments* (Bingley, UK: Emerald, Research in Experimental Economics, Volume 12, 2008).

Keeney, Ralph L., "Risk Independence and Multiattributed Utility Functions," *Econometrica*, 41(1), January 1973, 27-34.

Kreps, David M., and Porteus, Evan L., "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46(1), January 1978, 185-200.

Loewenstein, George; O'Donoghue, Ted, and Rabin, Matthew, "Projection Bias in Predicting Future Utility," *Quarterly Journal of Economics*, 118(4), November 2003, 1209-1248.

Rabin, Matthew, "Risk Aversion and Expected Utility Theory: A Calibration Theorem," *Econometrica*, 68, 2000, 1281-1292.

Richard, Scott F., "Multivariate Risk Aversion, Utility Independence and Separable Utility Functions," *Management Science*, 22(1), September 1975, 12–21.

Schmidt, Ulrich, "A Measurement of the Certainty Effect," *Journal of Mathematical Psychology*, 42, 1998, 32-47.

Wooldridge, Jeffrey M., *Econometric Analysis of Cross Section and Panel Data* (MIT Press: Cambridge, 2002).

Appendix: Additional Figures with Detailed Data (NOT FOR PUBLICATION)

Figure A1: Portfolio Allocations by Annual Interest Rate from Andreoni and Sprenger [2012a] Experiments With No Front End Delay

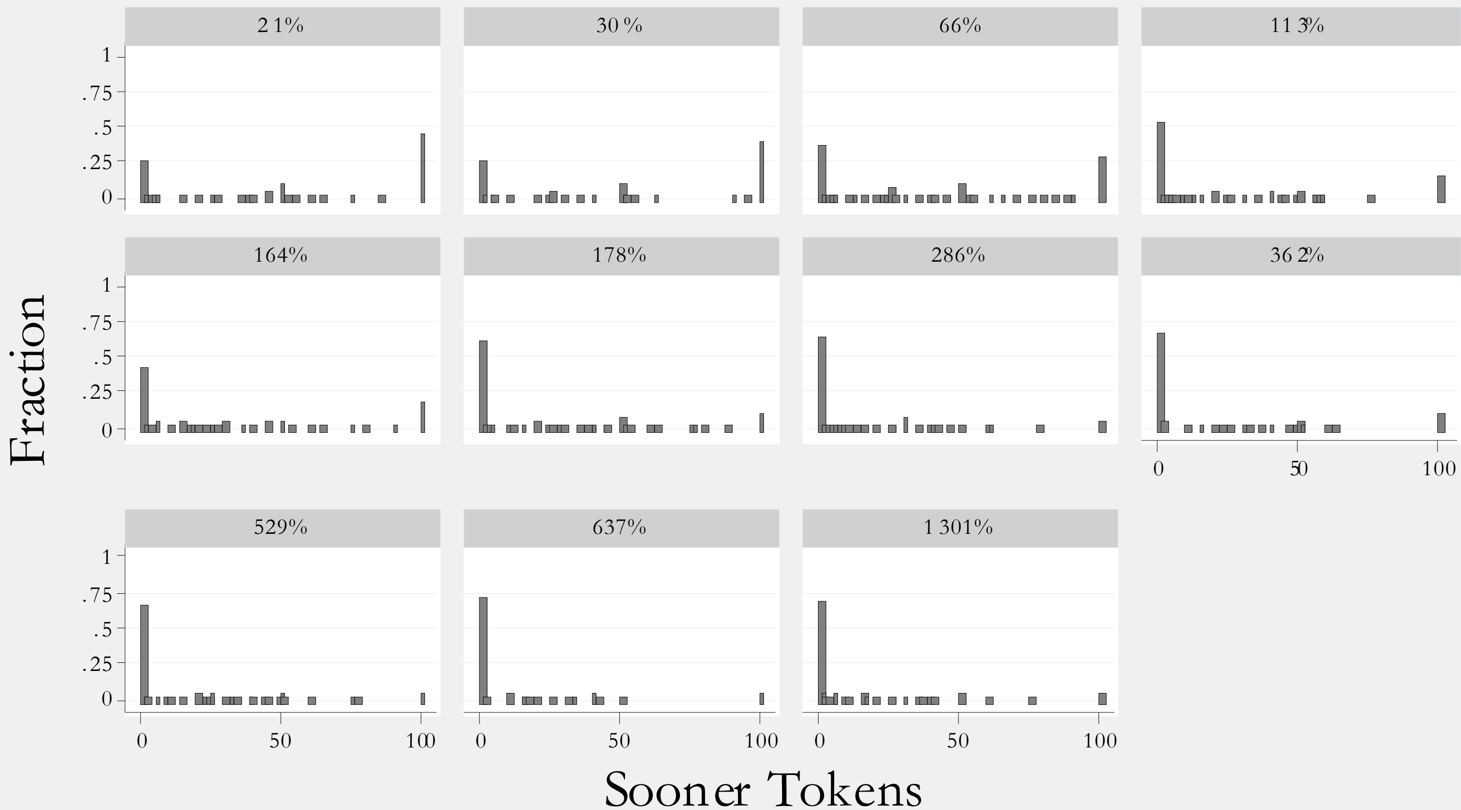

Figure A2: Portfolio Allocations by Annual Interest Rate from Andreoni and Sprenger [2012a] Experiments With Front End Delay of One Week

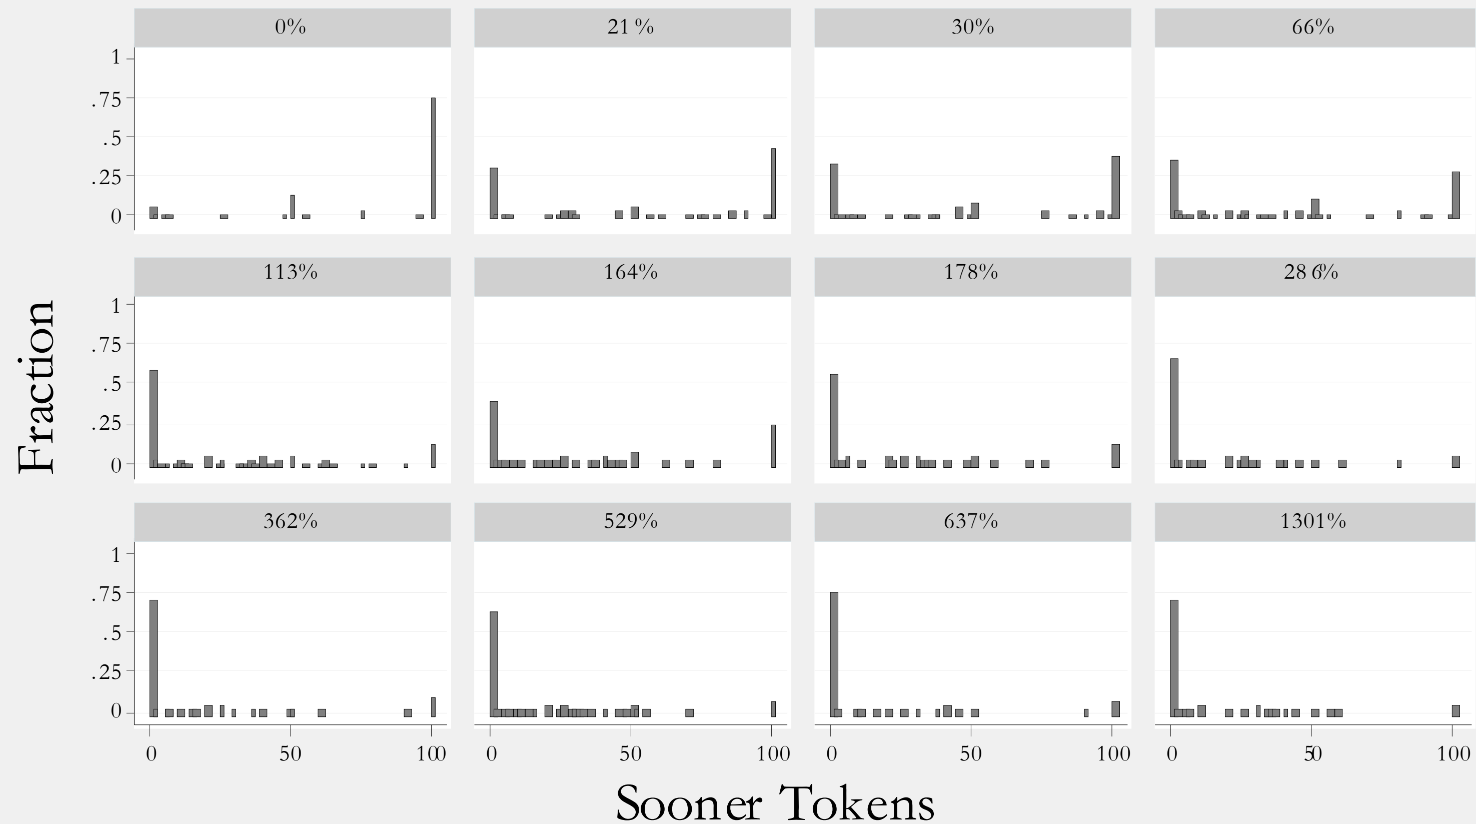

Figure A3: Portfolio Allocations by Annual Interest Rate from Andreoni and Sprenger [2012a] Experiments With Front End Delay of Five Weeks

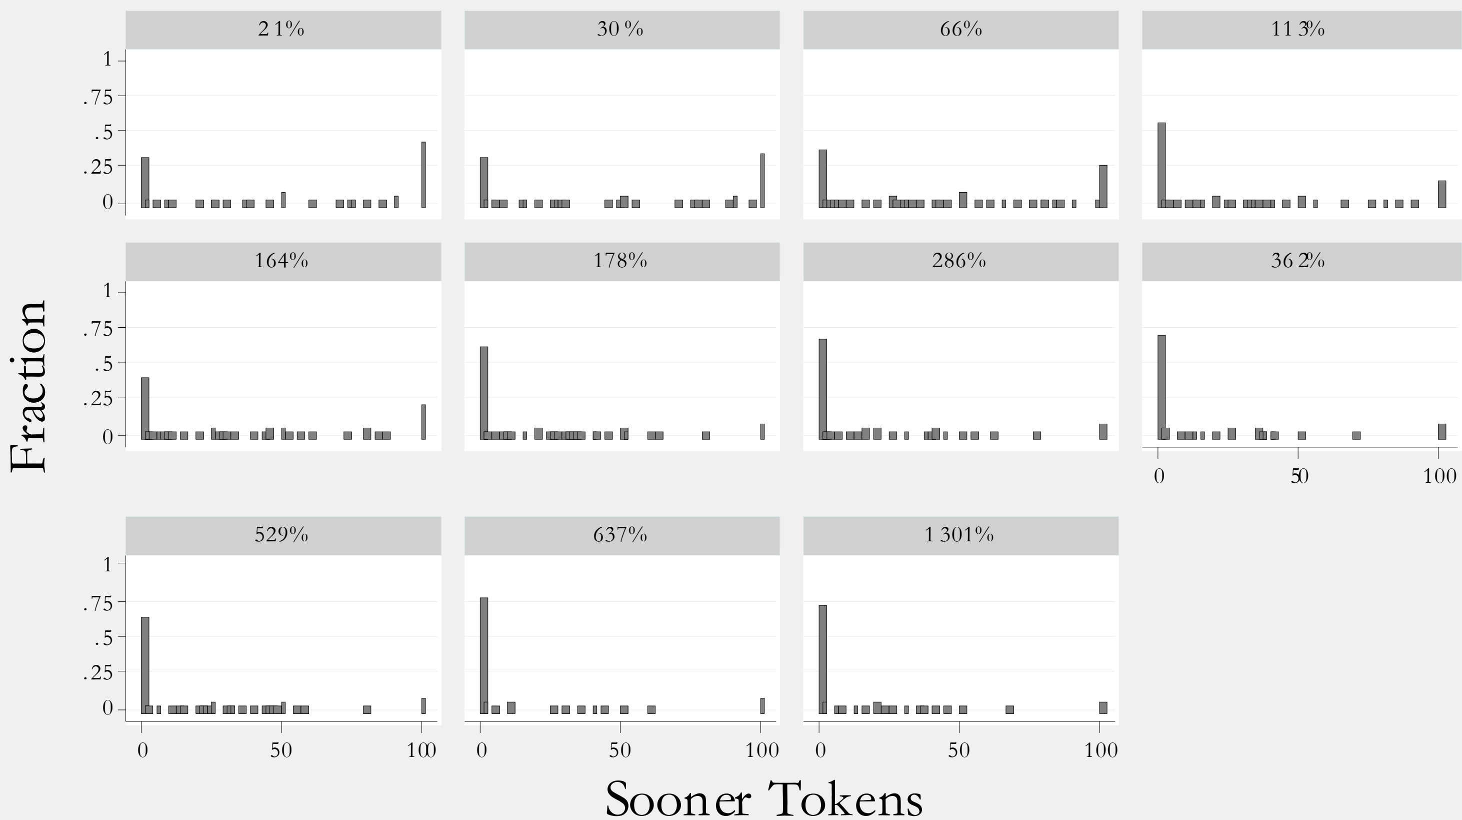

Figure A4: Portfolio Choices with Zero Interest Rates
from Andreoni and Sprenger [2012b] Experiments

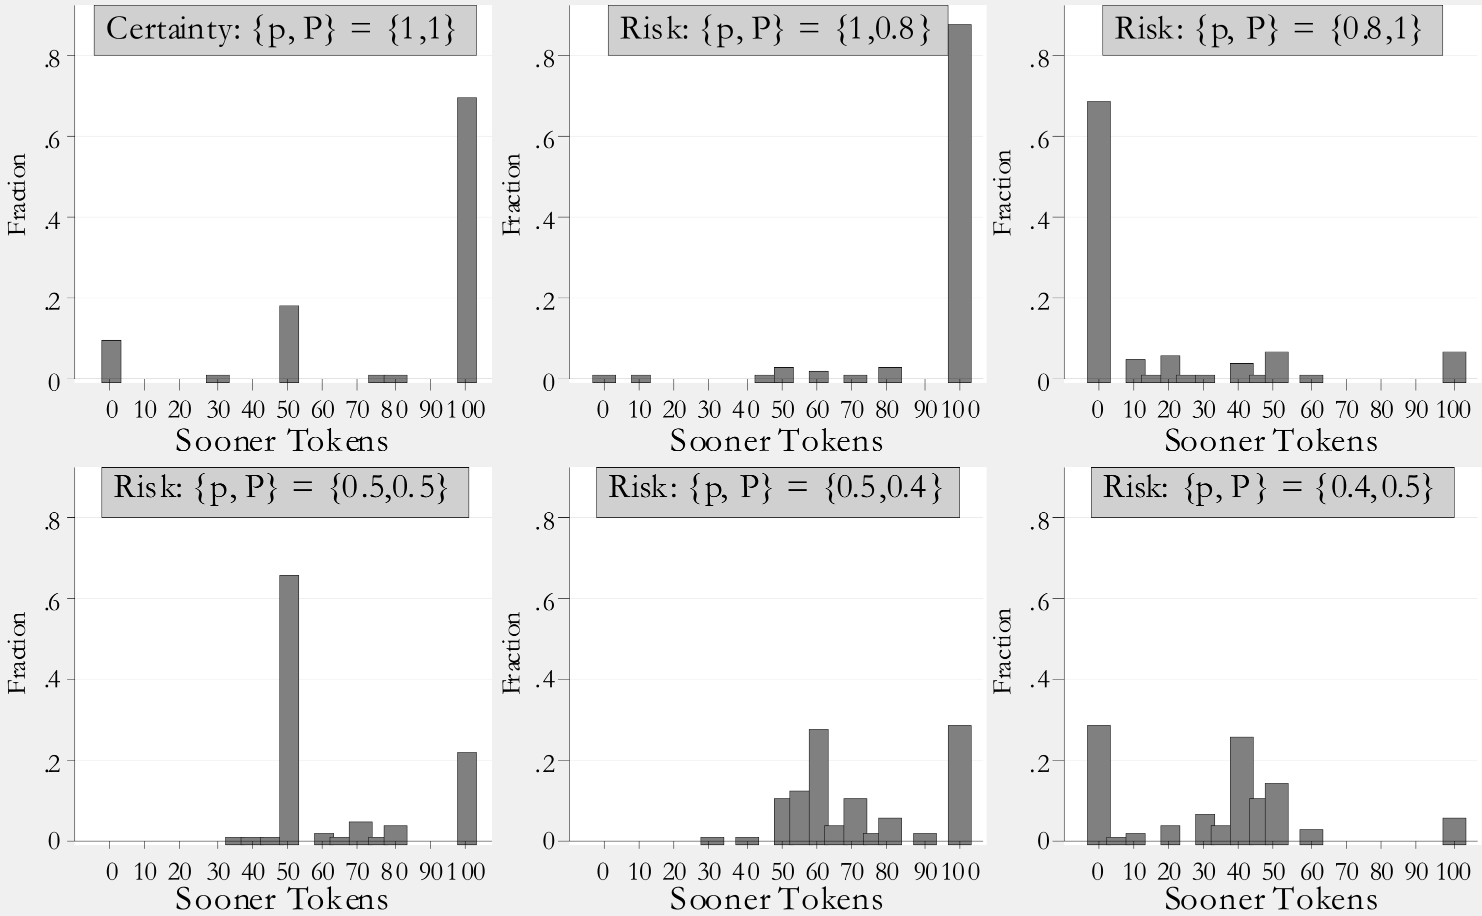

Figure A5: Portfolio Choices with Positive Interest Rates from Andreoni and Sprenger [2012b] Experiments