

Robustness of Inferences in Risk and Time Experiments to Lifecycle Asset Integration*

AJ A. Bostian¹ and Christoph Heinzl²

¹*School of Social Sciences and Humanities, University of Tampere, Finland*

²*INRAE SMART-LERECO, Agrocampus Ouest, Rennes, France*

December 1, 2020

Abstract

Participants in an experiment can engage in unobservable *asset integration*, mentally incorporating their own non-experimental “field” resources into an otherwise controlled scenario. This paper extends asset integration to include intertemporal tradeoffs like consumption smoothing. A model of “lifecycle asset integration” shows that exogenous and endogenous field resources cause different interference patterns. Exogenous resources cannot be affected by the experiment, and so their interference can be controlled by accounting for their level. Endogenous resources, by contrast, are highly substitutable with the experiment, and their interference can be controlled only by modeling the entire experiment-field interaction. The model’s practical implications are investigated in the context of three classic laboratory experiments on risk and time: one static (Holt and Laury, 2002) and two dynamic (Andersen et al., 2008; Andreoni and Sprenger, 2012). As interference worsens, decisions in these tasks tend to exhibit a kind of attenuation bias toward less risk aversion and more patience. Interference occurs reliably when field resources are on household scales, but amounts on the scale of pocket change can also cause problems.

Keywords: risk aversion, consumption smoothing, discount rate, asset integration, experiment

JEL classification: C90, D81, D15

*We thank Glenn Harrison, Charlie Holt, and Olivier l’Haridon for their helpful comments; as well as seminar participants at the University of Alabama, Australian National University, BETA Strasbourg, CEE-M Montpellier, CREM Rennes, INRA Rennes, New York University Abu Dhabi, Oregon State University, Portland State University, Resources for the Future, University of Technology Sydney, Umeå University, University of Virginia, ARIA (ASSA meeting), and EGRIE. We also thank Benoît Carré for very helpful research assistance. This project was partially supported by the European Union’s Seventh Framework Programme FP7/2007-2011, Grant Agreement no. 290693 (FOODSECURE), and by the J. William Fulbright Foundation.

1 Introduction

Experiments that elicit risk and time preferences are susceptible to *asset integration*, a design bypass that occurs when participants mentally incorporate their own financial resources into a supposedly controlled scenario. Because asset integration is triggered by the scenario itself, its impacts cannot be neutralized with randomization. As an unobservable mental process, it can introduce a particularly unfortunate confound into *post hoc* utility estimates.

Although asset integration was first discussed in a static context (Kahneman and Tversky, 1979), it can appear in dynamic settings just as easily. But, asset integration’s current adaptation to time is rudimentary: the non-experimental or “field” environment consists merely of recurring static resources. If this is the correct interpretation of asset integration, then the confound can be eliminated rather easily by including those resource levels in the utility argument.

However, participants almost assuredly do not regard their own field environments as mere successions of resources that they must take as given. Instead, they sew up those field resources into a coherent lifecycle plan. That process taps intertemporal tradeoffs that do not emerge under the recurring-static view. A well-known example is consumption smoothing, which distributes temporally-unbalanced resources more evenly across time. We use the term “lifecycle asset integration” to denote the complications that arise when lifecycle factors confound an experiment.

In this paper, we develop a model of lifecycle asset integration, and examine its implications for experiments that investigate risk and time preferences. This model is rooted in the two-period consumption-saving framework, the canonical theory of lifecycle decision making under risk (Drèze and Modigliani, 1966, 1972; Leland, 1968; Sandmo, 1970; Rothschild and Stiglitz, 1971; Kimball, 1990; Eeckhoudt and Schlesinger, 2008; Kimball and Weil, 2009). Our extension merges an experimenter’s controlled experimental incentives with a participant’s existing field environment, producing a joint optimization problem that involves experimental and field smoothing.

Our main theoretical result highlights a key difference in how exogenous and endogenous field resources interact with an experiment. Exogenous resources are very much like the recurring-static conception of the field. They have no associated smoothing instrument (such as saving), and so they cannot be moved across time. Endogenous resources, on the other hand, do involve a smoothing instrument, and therefore can be moved across time. Endogenous field instruments that would be relevant in this context include participants' credit cards and bank accounts.

An experiment has no effect whatsoever on exogenous field resources. As a consequence, the appropriate *post hoc* correction for exogenous field resources is indeed to control for their levels in the utility argument (Andersen et al., 2018). However, this simple *ceteris paribus* control strategy is inadequate for endogenous field resources, because experimental and field smoothing instruments can substitute for each other.

We argue that the marginal rate of substitution (MRS) between experimental and field smoothing is likely to be near 1 in most circumstances, implying near-perfect substitution. In other words, the existence of the experiment will actually alter the amount of field smoothing, another component of the utility argument. Simply conditioning on pre-existing field outcomes is insufficient in this case, because the experiment itself will change those outcomes. The appropriate correction here is much more onerous: modeling the entire experiment-field interaction, including the MRS.

To illustrate the practical effects of lifecycle asset integration, we numerically investigate the model's predictions in the context of three classic laboratory experiments: Holt and Laury (2002)'s risk aversion task (HL), Andersen et al. (2008)'s discounting task (AHLR), and Andreoni and Sprenger (2012)'s discounting task (AS). In each case, our first step is to cast the experiment's incentives into the notation of our model. This allows us to treat the experiment on its own terms first, and then add different integration assumptions. Specifically, we examine how the experimental predictions change as integration changes, *holding preferences constant*. If the predictions are substantively different, that experiment

is not robust to lifecycle asset integration. The “controlled” experimental observations can be easily contaminated by field interactions.

Because the dividing line between experimental and field incentives is very clear with laboratory experiments, those three designs are relatively easy to adapt to the model. Interestingly, asset integration does not spark uniform concern within the laboratory literature. At one extreme, asset integration’s relevance to static choice-bracketing experiments has been debated vigorously (Read et al., 1999; Rabin and Weizsäcker, 2009). At the other, many temporal experiments simply assume asset integration from the start (Coller and Williams, 1999; Cubitt and Read, 2007; Andersen et al., 2008, 2014; Andreoni and Sprenger, 2012). Critically, all experimental gains – even the gains from static tasks – are spent within participants’ lifecycle plans, making every experiment potentially vulnerable to lifecycle asset integration.

Our first example, the HL task, requires participants to choose between pairs of safe and risky lotteries. As a static decision, HL does not have an experimental smoothing instrument, and so it evades the substitution problem between field and experimental smoothing. It is still exposed to the other issues.

Omitting exogenous field resources situates *post hoc* analysis at the wrong background level. HL assumes that this level is \$0, but we find a strong sensitivity to that assumption. Adding just \$0.20 of exogenous field resources causes HL decisions to change. Increasing that level to just \$7 causes those decisions to be indistinguishable from risk neutrality.

Omitting endogenous field resources ignores how the task activates consumption smoothing. From the perspective of a participant’s lifecycle plan, HL payoffs are an instantaneous windfall that the participant will want to smooth across time via the field instrument. The larger the experimental windfall, the more field smoothing will occur. We find that the ability to smooth acts as a kind of self-insurance, allowing the participant to mitigate experimental risk in the field over time. This self-insurance pushes experimental decisions toward risk neutrality as well, even though preferences are calibrated as risk averse.

Our second example, the AHLR task, requires participants to choose between pairs of current and future payoffs. Although AHLR does not explicitly call its decision “saving,” we show that its decision problem nevertheless contains a latent variable corresponding to our model’s experimental smoothing variable. AHLR is exposed to all issues posed by lifecycle asset integration, but the omission of endogenous field resources results in two notable complications.

First, as with HL, a participant can potentially smooth an AHLR experimental windfall using a field instrument that the experimenter cannot observe. However, because the AHLR decision is itself a smoothing decision, the participant can also satisfy that smoothing desire during the experiment using the observable instrument. This fact actually underscores the bigger problem: participants now have two *highly substitutable* smoothing instruments at their disposal.

Intuitively, a participant will try to save as much as possible using the instrument that provides the better outcome. The meaning of “better” in this context depends on the field consumption path. If that path already has high future consumption, for example, offering large experimental returns may not induce any saving. This is quite different from the predictions that arise when the experimental smoothing instrument is assumed to operate by itself.

Our third example, the AS task, can be viewed as an extension of AHLR that allows participants to choose their own smoothing amounts. That freedom of choice very clearly illustrates the MRS between experimental and field smoothing. In particular, it is quite easy to drive experimental saving to its boundary values by making the experimental incentives too stingy or too rich relative to the field. Andreoni and Sprenger do, in fact, observe an unusually large fraction of boundary decisions in their data.

HL, AHLR, and AS each estimate deep structural parameters under expected utility (EU). Our model uses recursive utility (RU) instead (Kimball and Weil, 2009). The main reason is flexibility: RU allows smoothing preferences to be totally uncoupled from risk

preferences, while EU fuses them. But, because EU is a special case of RU, our framework covers all EU-based analysis. To be clear, our main theoretical result does not hinge on whether a participant is an EU or RU decision maker. Rather, when asset integration has a *lifecycle* nature, it is easier to illustrate the interference starting from the position that relative risk aversion (RRA) and the elasticity of intertemporal substitution (EIS) are not functionally bound to each other. That stance allows us to calibrate each domain with plausible values taken from its own literature.

Prior research hints at aspects of our results. Cubitt and Read (2007) discuss interference between temporal experimental choices and field variables under EU, but risk does not enter their assessment. Schechter (2007) performs a similar evaluation, calibrating an intertemporal utility function with the results of a static risk task. Coble and Lusk (2010) examine AHLR with isoelastic RU preferences, as do Miao and Zhong (2015) when examining AS. Both find empirical support for RU over EU, but neither include any asset integration. Administering AS to a large Danish cohort, Epper et al. (2020) find that patience positively correlates with wealth across the lifecycle. Below, we gather all these earlier concerns about the experiment-field interaction, along with some new ones, into a single analytical framework.

2 Lifecycle Asset Integration

To formally capture the interaction between an experimental stimulus and a participant’s lifecycle path, we extend a two-period consumption-saving model.¹ The purpose of this extension is to differentiate field incentives from experimental ones.

To that end, field elements of the model are denoted with f superscripts, and experimental

¹The two-period framework permits only two temporal actions: smoothing forward in time, or backward in time. We could certainly obtain a more nuanced time path with a multiperiod value function. But, a model of that sort would generate essentially the same first-order conditions as ours, while greatly increasing the difficulty of the numerical exercises (solving for policy functions instead of scalars). Moreover, we could not appeal to the theoretical corpus for intuition, because that literature largely operates with two periods. Because we are not concerned with the time path of field smoothing *per se*, but with whether that smoothing interferes with experiments, we stick with this simpler “forward or backward” approach.

elements with e superscripts. Risky variables are denoted with tildes. By convention, risk occurs in the second period. Field and experimental risks are considered orthogonal by construction.

The field incentives are the first-period income y_1^f , the second-period exogenous risky income \tilde{y}_2^f , and the second-period gross return to saving R_2^f (net return r_2^f). The experimental incentives are analogously y_1^e , \tilde{y}_2^e , and R_2^e .² Intuitively, the field incentives are outside the experimenter’s control, while the experimental incentives are the experimenter’s own manipulation.

During the first period, the participant chooses field saving s_1^f and experimental saving s_1^e to maximize the RU objective over lifecycle consumption c_1^f, \tilde{c}_2^f :

$$\max_{s_1^f, s_1^e} u(c_1^f) + \beta u(CE(\tilde{c}_2^f)) \quad \text{s.t.} \quad \begin{cases} y_1^f + y_1^e = c_1^f + s_1^f + s_1^e \\ \tilde{y}_2^f + \tilde{y}_2^e + s_1^f R_2^f + s_1^e R_2^e = \tilde{c}_2^f \\ -\tilde{y}_2^f \leq s_1^f \leq y_1^f \\ 0 \leq s_1^e \leq y_1^e \end{cases} \quad (1)$$

We call s_1^f and s_1^e “saving,” as does most of the literature. That terminology certainly conveys s_1^f ’s and s_1^e ’s function in (1). But, those variables more precisely reflect a generic two-channel smoothing framework that allocates consumption across time. This nuance allows us to bring experiments under the umbrella of (1) that do not explicitly invoke the language of saving, but nevertheless set up a latent smoothing instrument that behaves like s_1^e .³

RU’s notion of time preference has two ingredients: the utility discount factor β , and the

²Limiting risk to the exogenous channels \tilde{y}_2^f and \tilde{y}_2^e greatly simplifies our discussion of risk attitudes and risk responses. When risk exposure is endogenous with a choice variable, the lifecycle risk response becomes quite complex (Eeckhoudt and Schlesinger, 2008). Risks on \tilde{R}_2^f and \tilde{R}_2^e would be prime examples of endogenous exposure, because the amounts at risk $s_1^f \tilde{R}_2^f$ and $s_1^e \tilde{R}_2^e$ would depend on the choice variables s_1^f and s_1^e . Fortunately, none of our three examples involve manipulations of return risk, and so we do not lose much intuition by narrowing our focus in this way.

³Only a few experiments intentionally style their tasks as “saving” (Ballinger et al., 2003, 2011; Brown et al., 2009; Filiz-Ozbay et al., 2015).

intertemporal felicity function u that controls consumption smoothing. This is an important difference from EU, where time preference is considered to be simply β or some variant. The risk preference ψ determines the certainty equivalent of future consumption

$$CE(\tilde{c}_2^f) \equiv \psi^{-1}\left(E_1^f E_1^e \left[\psi\left(\tilde{c}_2^f\right)\right]\right)$$

Unlike u , ψ is an EU function. RU thus contains both a utility-of-wealth function ψ , and a utility-of-consumption function u . The special case $u = \psi$ collapses RU to EU (Kreps and Porteus, 1978).

In lifecycle models like this one, nearly any incentive – field or experimental – will activate β , u , and ψ simultaneously (Gollier, 2001). That is equally true of risk (Kimball and Weil, 2009), somewhat counterintuitively. The lifecycle risk response differs from the more familiar static risk response in two additional ways. First, a risk's n^{th} moment activates the $n + 1^{\text{th}}$ utility derivative, not the n^{th} (Eeckhoudt and Schlesinger, 2008). The lifecycle risk response therefore depends on ψ derivatives *higher* than ψ'' . By corollary, the well-known Arrow-Pratt coefficient $-\psi''/\psi'$ has nothing to say about a participant's reaction to lifecycle risk.

The first two constraints in (1) show how experimental and field assets become integrated into the participant's lifecycle plan. The vehicle is lifecycle consumption. Both first-period consumption $c_1^f = (y_1^f + y_1^e) - (s_1^f + s_1^e)$ and second-period consumption $\tilde{c}_2^f = (\tilde{y}_2^f + \tilde{y}_2^e) + (s_1^f R_2^f + s_1^e R_2^e)$ contain a mixture of experimental and field resources.

The properties of u and ψ that ensure (1) has a unique maximum also guarantee interior equilibrium consumption levels c_1^{f*} and c_2^{f*} . For that reason, we simply assert the equality of the first two constraints. However, the fact that the consumption path is interior does not mean that the saving decisions are as well. Without additional information about preferences and incentives, the constraints on s_1^{f*} and s_1^{e*} must stay as inequalities.

Field saving s_1^f can be positive or negative. We place a rather loose restriction on saving and borrowing via this channel: the participant's own lifecycle field income. We place

a much stronger restriction on experimental saving s_1^e : the participant cannot borrow at all during the experiment, or save more than y_1^e . This reflects a common design constraint. For ethical reasons, participants are typically prohibited from investing – and potentially losing – their own field resources in an experiment. Their experimental decisions must always fit within the resource levels endowed by the experimenter.

The interpretation of the field resources y_1^f and \tilde{y}_2^f is context-specific. As a rule, those terms express incentives that the participant *perceives* to be relevant to the experimental decision s_1^e , but are not actually part of the experiment. To reiterate, these perceptions are not entirely under the experimenter’s control, nor are they fully observable. Thus, y_1^f and \tilde{y}_2^f should not be read merely as attributes that the experimenter could account for in principle. They can also reflect traits that the experimenter has no hope of observing.

Even though the experimental incentives are controlled, the participant’s experimental and field decisions remain tightly coupled. The nature of that coupling can be seen by writing (1)’s first-order conditions in Euler form:⁴

$$s_1^f : E_1^f E_1^e \left[\beta \cdot \left(\frac{u'(CE(\tilde{c}_2^f))}{u'(c_1^f)} / \frac{\psi'(CE(\tilde{c}_2^f))}{\psi'(\tilde{c}_2^f)} \right) \cdot R_2^f \right] - 1 \stackrel{\leq}{\geq} 0 \quad (2a)$$

$$s_1^e : E_1^f E_1^e \left[\beta \cdot \left(\frac{u'(CE(\tilde{c}_2^f))}{u'(c_1^f)} / \frac{\psi'(CE(\tilde{c}_2^f))}{\psi'(\tilde{c}_2^f)} \right) \cdot R_2^e \right] - 1 \stackrel{\leq}{\geq} 0 \quad (2b)$$

Importantly, the first Euler equation remains pertinent even if the experimental task is static (i.e., $s_1^e \equiv 0$), because static experimental incentives can still alter c_1^f and \tilde{c}_2^f (and hence s_1^f).

Individually, (2a) and (2b) are examples of a discounted-return equilibrium condition $E_t(m_{t+1}R_{t+1}) \stackrel{\leq}{\geq} 1$ that arises ubiquitously in dynamic models (Cochrane, 2005). Like that generic condition, the left sides here take the form of discounted expected returns. The quantity corresponding to the discount function m_{t+1} is the participant’s stochastic discount

⁴These conditions contain inequalities because s_1^{f*} and s_1^{e*} are not necessarily interior.

factor (SDF)

$$\beta \cdot \left(\frac{u' (CE (\tilde{c}_2^f))}{u' (c_1^f)} / \frac{\psi' (CE (\tilde{c}_2^f))}{\psi' (\tilde{c}_2^f)} \right) \quad (3)$$

As a system, (2a) and (2b) are akin to the first-order conditions of a multi-asset portfolio model (Ingersoll, 1987). A key feature of that setting is the fact that all portfolio allocations are pinned down by the same m_{t+1} . Our model shares this characteristic: the same SDF (3) determines both s_1^f and s_1^e . This carries an important equilibrium implication: a change to either one of the saving amounts will alter the SDF, and thereby affect the other's discounted return – and hence change the other saving amount as well.

In light of that observation, the participant's SDF (3) can be best described as the behavioral transmission mechanism between the experiment and the field. This transmission is regulated by β and two marginal rates of substitution. The one involving u captures the consumption-smoothing implications of the joint decision, and the one involving ψ captures the risk-aversion implications.

The fact that the SDF contains all three behavioral primitives has an important ramification for interpreting the experimental outcome s_1^e . Namely, the presence of lifecycle asset integration will cause *all preference dimensions* – consumption smoothing, risk aversion, and discounting – to activate simultaneously in response to any set of incentives. Hence, even if the experimenter's *intent* is to design a manipulation $\{y_1^e, \tilde{y}_2^e, R_2^e\}$ that activates only risk attitudes, or only smoothing attitudes, or only the pure rate of time preference, lifecycle asset integration will nevertheless activate everything.

The participant's elasticity of substitution between experimental and field decisions makes

the practical implications of the SDF transmission quite clear:⁵

$$\begin{aligned} \epsilon^{e,f} &= \frac{ds_1^f}{ds_1^e} \cdot \frac{s_1^e}{s_1^f} \\ &= - \frac{u''(c_1^f) + \beta \left[u''(\tilde{c}_2^f) CE'(\tilde{c}_2^f)^2 + u'(\tilde{c}_2^f) CE''(\tilde{c}_2^f) \right] \cdot \frac{1}{2} (R_2^f + R_2^e) R_2^e}{u''(c_1^f) + \beta \left[u''(\tilde{c}_2^f) CE'(\tilde{c}_2^f)^2 + u'(\tilde{c}_2^f) CE''(\tilde{c}_2^f) \right] \cdot \frac{1}{2} (R_2^f + R_2^e) R_2^f} \cdot \frac{s_1^e}{s_1^f} \end{aligned} \quad (4)$$

The MRS $\left| ds_1^f / ds_1^e \right|$ describes how well the participant's field saving can replace experimental saving, and vice versa. Ideally, this quantity would be 0. The experimental and field decisions would not affect each other at all in that case, signaling that the experimenter's manipulation is truly exogenous.

However, the numerator and denominator of the MRS differ by only their very last terms, R_2^e and R_2^f . Thus, barring any implausibly extreme experimental incentives, the MRS is likely to be close to 1. This unfortunately means that experimental smoothing and field smoothing can perfectly substitute for each other. In that case, the experimenter must worry about field contamination from not only the exogenous sources y_1^f and \tilde{y}_2^f , but also the endogenous source $s_1^f R_2^f$. The participant's own choices during the experiment fortunately cannot affect the former. But, per the MRS, they can affect the latter.

Finally, because many experimental studies assume EU, it is worth noting the EU SDF's behavior when it is viewed as an RU special case. As a rule, EU requires the "reduction of compound lotteries" axiom to hold in all circumstances. In lifecycle settings, this means that the axiom must hold both within and across time. RU loosens that requirement into "temporal consistency," which requires conformity within time only (Kreps and Porteus, 1978; Selden, 1978). Temporal consistency materializes in the SDF (3) as the distinction between risk substitution and intertemporal substitution.

As u and ψ become more similar, that distinction becomes irrelevant. This is reflected

⁵Appendix A contains the derivation of $\epsilon^{e,f}$.

in the SDF, which collapses to

$$\beta \cdot \frac{u'(\tilde{c}_2^f)}{u'(c_1^f)} = \beta \cdot \frac{\psi'(\tilde{c}_2^f)}{\psi'(c_1^f)}$$

in the EU special case $u = \psi$. In this circumstance, it notably does not matter whether the felicity function is taken to be an “intertemporal preference” u or a “risk preference” ψ . The same decisions will be made under either interpretation.

That result, if valid, provides a powerful design shortcut. The experimenter can elicit a participant’s risk preference and immediately treat it as the intertemporal preference, or vice versa. Problematically, a good deal of empirical literature, particularly from macroeconomics, rejects the hypothesis that risk substitution and intertemporal substitution have the same elasticity (Epstein and Zin, 1989, 1991; Bansal and Yaron, 2004).

3 Merging Lifecycle Asset Integration into Three Experiments

3.1 Holt and Laury

As a static experiment, HL has no internal concept of time. Even so, we will motivate it with temporal notation consistent with (1). Within the context of HL itself, that notation is pure surplusage: nothing would be gained or lost by adding or removing it. But, including it from the outset makes the transition to our model much easier.

Table 1 presents the baseline HL task, a multiple price list (MPL) of ten safe and risky lotteries. The payoffs for each safe lottery $\tilde{y}_2^{e, safe}$ and each risky lottery $\tilde{y}_2^{e, risky}$ remain the same throughout the MPL. The payoffs of $\tilde{y}_2^{e, risky}$ always have more spread than those of $\tilde{y}_2^{e, safe}$.

Moving down the MPL, the probabilities increasingly favor each lottery’s high payoff.

Line	$\tilde{y}_2^{e,safe}$ (\$)		$\tilde{y}_2^{e,risky}$ (\$)	
1	1/10 of 2.00	9/10 of 1.60	1/10 of 3.85	9/10 of 0.10
2	2/10 of 2.00	8/10 of 1.60	2/10 of 3.85	8/10 of 0.10
3	3/10 of 2.00	7/10 of 1.60	3/10 of 3.85	7/10 of 0.10
4	4/10 of 2.00	6/10 of 1.60	4/10 of 3.85	6/10 of 0.10
5	5/10 of 2.00	5/10 of 1.60	5/10 of 3.85	5/10 of 0.10
6	6/10 of 2.00	4/10 of 1.60	6/10 of 3.85	4/10 of 0.10
7	7/10 of 2.00	3/10 of 1.60	7/10 of 3.85	3/10 of 0.10
8	8/10 of 2.00	2/10 of 1.60	8/10 of 3.85	2/10 of 0.10
9	9/10 of 2.00	1/10 of 1.60	9/10 of 3.85	1/10 of 0.10
10	10/10 of 2.00	0/10 of 1.60	10/10 of 3.85	0/10 of 0.10

Table 1: Holt and Laury’s baseline MPL

Lifecycle Variable	HL “Safe”	HL “Risky”
y_1^e	0	0
s_1^e	-	-
\tilde{y}_2^e	$\tilde{y}_2^{e,safe}$	$\tilde{y}_2^{e,risky}$
R_2^e	-	-
y_1^f	-	-
\tilde{y}_2^f	-	-
s_1^f	-	-
R_2^f	-	-

Table 2: Translation between model (1) and HL

The two lotteries have equal means at line 5. The safe lottery’s mean is higher before that line, and the risky lottery’s mean is higher after that line.

On each line, the participant indicates a preference for the safe or risky lottery. That decision is governed by the comparison

$$E_1^e \left[\psi \left(\tilde{y}_2^{e,safe} \right) \right] \gtrless E_1^e \left[\psi \left(\tilde{y}_2^{e,risky} \right) \right] \quad (5)$$

Given the ordering of the MPL, a risk-neutral participant would switch from safe to risky for good at line 5. A risk-averse participant would switch later.

Table 2 summarizes how these incentives translate to model (1).⁶ Two items are of

⁶Because the first period is usually taken to be the “present” and the second period the “future,” a more natural timeline might place the HL experimental resources in the first period. However, that would require

particular note. First, because HL does not have a smoothing instrument, the s_1^e aspect does not apply. Second, HL does not address the field at all.

The comparison analogous to (5) under lifecycle asset integration is

$$u\left(c_1^{f,safe}\right) + \beta u\left(CE\left(\tilde{c}_2^{f,safe}\right)\right) \gtrless u\left(c_1^{f,risky}\right) + \beta u\left(CE\left(\tilde{c}_2^{f,risky}\right)\right)$$

Applying simplifications from Table 2 and expanding terms yields

$$\begin{aligned} & u\left(y_1^f - s_1^{f,safe}\right) + \beta u\left(\psi^{-1}\left(E_1^f E_1^{e,safe}\left[\psi\left(\tilde{y}_2^f + \tilde{y}_2^{e,safe} + s_1^{f,safe} R_2^f\right)\right]\right)\right) \gtrless \quad (6) \\ & u\left(y_1^f - s_1^{f,risky}\right) + \beta u\left(\psi^{-1}\left(E_1^f E_1^{e,risky}\left[\psi\left(\tilde{y}_2^f + \tilde{y}_2^{e,risky} + s_1^{f,risky} R_2^f\right)\right]\right)\right) \end{aligned}$$

Critically, the participant's field incentives $\left\{y_1^f, \tilde{y}_2^f, R_2^f\right\}$ do not change on either side of this comparison. The question is whether the participant's field decision s_1^f does.

If s_1^f does not actually change, then the lifecycle comparison (6) collapses right back to the static one. To see this, note that after simplifying and rearranging terms, the comparison becomes

$$E_1^f E_1^{e,safe}\left[\psi\left(\tilde{y}_2^{e,safe} + \left(\tilde{y}_2^f + s_1^f R_2^f\right)\right)\right] \gtrless E_1^f E_1^{e,risky}\left[\psi\left(\tilde{y}_2^{e,risky} + \left(\tilde{y}_2^f + s_1^f R_2^f\right)\right)\right]$$

when $s_1^{f,safe} = s_1^{f,risky}$. This expression is nothing more than (5) with an additional field term $\tilde{\omega}_2^f = \tilde{y}_2^f + s_1^f R_2^f$ in the utility argument. If this is the correct assumption about s_1^f 's behavior, then the only essential refinement to HL is to include the participant's background field assets and risks (Heinemann, 2008; Harrison et al., 2017; Andersen et al., 2018).

However, it is more likely that each side of the comparison will yield different values of s_1^f . The reason is that the total risk \tilde{c}_2^f has different means and variances on each side. Those either breaking (1)'s timing convention (where risk can fall only in the second period), or adding more periods. The first option would obfuscate otherwise crisp theoretical predictions about the consumption-smoothing and precautionary motives, while the second would result in more convoluted comparisons. Both would add complexity that is not needed to see the intuition on HL interference from lifecycle factors. To the extent there is an ideal way to frame the timing, the fact that our comparison (6) reduces to HL's (5) after removing all field resources indicates our version respects both HL and the goals of this exercise.

two moments

$$\begin{aligned} E_1(\tilde{c}_2^f) &= E_1^f(\tilde{y}_2^f) + E_1^e(\tilde{y}_2^e) + s_1^f R_2^f \\ V_1(\tilde{c}_2^f) &= V_1^f(\tilde{y}_2^f) + V_1^e(\tilde{y}_2^e) \end{aligned}$$

influence saving in well-known ways (Kimball, 1990; Eeckhoudt and Schlesinger, 2008; Bostian and Heinzl, 2018). The most consequential response is the one to $E_1(\tilde{c}_2^f)$, which triggers consumption smoothing. The second-most consequential is the one to $V_1(\tilde{c}_2^f)$, which triggers precaution.⁷

Because $E_1^{e, safe}(\tilde{y}_2^{e, safe}) \leq E_1^{e, risky}(\tilde{y}_2^{e, risky})$, the field saving amounts $s_1^{f, safe}$ and $s_1^{f, risky}$ will generally have different consumption-smoothing components. The lottery with the higher mean will induce more smoothing. Because the risk is dated the second period, that lottery will have lower s_1^f , and thus higher c_1^f .

Similarly, because $V_1^{e, safe}(\tilde{y}_2^{e, safe}) < V_1^{e, risky}(\tilde{y}_2^{e, risky})$, the field saving amounts $s_1^{f, safe}$ and $s_1^{f, risky}$ will have different precautionary components as well. The risky lottery will induce more precaution for sure. Because the risk is dated the second period, the risky lottery will have higher $s_1^{f, risky}$, and thus lower $c_1^{f, risky}$.

Comparing (5) to (6) reveals two specification errors that will arise during structural estimation by failing to include lifecycle asset integration. The root of both problems is the omitted field term $\tilde{\omega}_2^f = \tilde{y}_2^f + s_1^f R_2^f$.

The first is failing to control for the exogenous field resources \tilde{y}_2^f . The error here arises from implicitly assuming $\tilde{y}_2^f = 0$, thereby situating the participant's decisions at the wrong background level. Because HL's ψ specification allows increasing, decreasing, and constant RRA, centering the decision making at the correct \tilde{y}_2^f level is of paramount concern.

The second is failing to account for the endogenous field resources $s_1^f R_2^f$. This also injects a background-level problem, but that is not the only one. The compounded error arises from

⁷All of \tilde{c}_2^f 's higher moments also trigger precaution, but their contributions are negligible here.

Line	x_1^e (DKK)	x_2^e (DKK)	APR (%)	r_2^e (%)
1	3000	3075	5	2.5
2	3000	3152	10	5.1
3	3000	3229	15	7.6
4	3000	3308	20	10.3
5	3000	3387	25	12.9
6	3000	3467	30	15.6
7	3000	3548	35	18.3
8	3000	3630	40	21.0
9	3000	3713	45	23.8
10	3000	3797	50	26.6

Table 3: Andersen et al.’s MPL for a six-month delay

implicitly assuming that the participant’s field saving remains the same (at $s_1^f = 0$) on both sides of (6). This treats an endogenous field resource as if it is exogenous.

Casting endogenous resources as exogenous ones fails to appreciate the full array of preferences operating on \tilde{c}_2^f . The specific problem here is that the $s_1^f R_2^f$ component of \tilde{c}_2^f is governed by u , not just ψ . A structural model consisting of ψ alone would thus improperly assign all of the experimental decision to risk attitudes, when some of it is actually prompted by smoothing attitudes. The resulting “ ψ ” estimates would be an uninterpretable mash of risk and intertemporal preferences.

Without question, the second error requires a much more invasive correction than the first. The solution to omitting exogenous field resources is the usual prescription for omitted-variable bias: include those resources. Importantly, that correction does not require changing the structural model (5). The solution to omitting endogenous field resources, on the other hand, requires specifying how those resources interact with the experiment. Because that interaction pulls in field smoothing, the structural model itself must change to (6).

3.2 Andersen et al.

AHLR simultaneously elicits utility discount rates and utility curvature. The discounting task is a choice between two rewards x_1^e and x_2^e spaced τ days apart.

Lifecycle Variable	AHLR “Early”	AHLR “Late”
y_1^e	x_1^e	x_1^e
s_1^e	0	x_1^e
\tilde{y}_2^e	0	0
R_2^e	x_2^e/x_1^e	x_2^e/x_1^e
y_1^f	ω_1^f	ω_1^f
\tilde{y}_2^f	ω_2^f	ω_2^f
s_1^f	-	-
R_2^f	-	-

Table 4: Translation between model (1) and AHLR

Table 3 presents the MPL for $\tau = 180$ days. On each line, the earlier payment x_1^e is DKK 3,000, and the later payment x_2^e is higher than DKK 3,000. Moving down the MPL, x_2^e rises. With an analogy to saving in mind, the design requires the percentage increase from x_1^e to x_2^e to be larger than any conceivable field interest rate r_2^f over the same τ interval.

On each line, the participant indicates a preference for the earlier or later option. That decision is governed by the comparison

$$u\left(x_1^e + \omega_1^f\right) + \left(\frac{1}{1+\delta}\right)^\tau u\left(\omega_2^f\right) \gtrless u\left(\omega_1^f\right) + \left(\frac{1}{1+\delta}\right)^\tau u\left(x_2^e + \omega_2^f\right) \quad (7)$$

where δ is the utility discount rate, and ω_1^f and ω_2^f are field resources. Andersen et al.’s version of (7) also includes a breakdown of how long the participant draws out the consumption of x_1^e and x_2^e . Because that detail is tangential to our interference question, we focus on a special case where consumption is immediate in both periods.

As a temporal task, AHLR is relatively easy to adapt to (1). First, we can set the discounting parameter to $\beta = \left(\frac{1}{1+\delta}\right)^\tau$. Next, we can continue the saving analogy: instead of taking x_1^e during the first period, a participant can defer that amount and take x_2^e in the second. We refer to these as the “early” and “late” options.

Table 4 summarizes how these incentives translate to model (1). The first option is to take the earlier payment $y_1^e = x_1^e$, thus saving $s_1^{e,early} = 0$ and earning nothing later. The

second option is to save $s_1^{e,late} = y_1^e$, thus taking nothing early, but earning

$$x_2^e = s_1^{e,late} \cdot \frac{x_2^e}{x_1^e} \quad \rightarrow \quad x_2^e = s_1^{e,late} R_2^e$$

later. The quantity $R_2^e = x_2^e/x_1^e$ is the gross increase in the payment, exactly our notion of return. In both cases, $\tilde{y}_2^e = 0$.

Applying those notational changes to (7) yields the comparison

$$\begin{aligned} u\left(\left(y_1^e - s_1^{e,early}\right) + \omega_1^f\right) + \beta u\left(s_1^{e,early} R_2^e + \omega_2^f\right) &\geq \\ u\left(\left(y_1^e - s_1^{e,late}\right) + \omega_1^f\right) + \beta u\left(s_1^{e,late} R_2^e + \omega_2^f\right) &\leq \end{aligned} \quad (8)$$

We have intentionally left alone the obvious simplifications in this expression. The reason is that (8) clearly draws out the latent smoothing variable s_1^e , which has a direct analog in (1). Thus, AHLR's core comparison does indeed involve smoothing, even though it restricts the smoothing options to the extremes $s_1^e \in \{0, y_1^e\}$.

This task therefore activates the participant's intertemporal preference u and discounting β . To provide additional utility variation outside β 's influence, AHLR also includes a HL task. That identification is problematic from an RU perspective, because HL's riskiness activates ψ . This extra HL data generates the intended supplemental variation only when $u = \psi$ – the EU special case.

AHLR structural estimates do indeed assume EU. If that is the correct framework, then both tasks identify the same $u = \psi$ preference. If not, then the restriction $u = \psi$ results in estimates that mash together risk and intertemporal attitudes.

The final quantities to reconcile are the field assets. Comparing (8) with (1), the models match exactly when $\omega_1^f = y_1^f - s_1^f$ and $\tilde{\omega}_2^f = \tilde{y}_2^f + s_1^f R_2^f$. But, because AHLR does not include any notion of field smoothing or field risk, its internal conception of field assets is simply $\omega_1^f = y_1^f$ and $\tilde{\omega}_2^f = y_2^f$. That is, only the exogenous field assets are relevant.

Indeed, Andersen et al. describe ω_1^f and ω_2^f as “the optimized consumption stream based

on wealth and income that is perfectly anticipated before allowing for the effects of the money offered in the experimental tasks.” In other words, the field assets are frozen in place before the experiment starts, and the experiment cannot subsequently affect them. This begs the question of endogenous resources altogether.

The elasticity between field and experimental decisions underscores the problem with that interpretation. If the experiment truly cannot affect the field, this elasticity is 0. But, that outcome is difficult to square with (4): field smoothing is likely to be highly elastic with experimental smoothing if any lifecycle asset integration is present.

3.3 Andreoni and Sprenger

AS elicits utility discount rates using a task that is more open-ended than an MPL. The participant must split a money budget m^e into current and future payoffs using tokens t_1^e and t_2^e . This split takes place along the constraint

$$m^e = p_1^e t_1^e + p_2^e t_2^e$$

under token prices p_1^e and p_2^e .

Unusually, AS’s structural model has Stone-Geary utility:

$$\max_{t_1^e, t_2^e} u(p_1^e t_1^e - \omega_1^f) + \gamma \eta^\tau u(p_2^e t_2^e - \omega_2^f) \quad \text{s.t.} \quad m^e = p_1^e t_1^e + p_2^e t_2^e \quad (9)$$

The field resources ω_1^f and ω_2^f are interpreted as a minimum amount of background consumption that the participant must acquire in each period. The parameter γ reflects present bias, and η the daily utility discount factor. As with AHLR, this last aspect can be reconciled with (1) by setting $\beta = \gamma \eta^\tau$.

Reconciling the rest requires converting token units to saving units (money). Table 5 summarizes the translation to model (1). Normalizing the constraint in (9) by p_1^e reveals an

Lifecycle Variable	AS Variable
y_1^e	m^e
s_1^e	$p_2^e t_2^e$
\tilde{y}_2^e	0
R_2^e	1
y_1^f	$-\omega_1^f$
\tilde{y}_2^f	$-\omega_2^f$
s_1^f	-
R_2^f	-

Table 5: Translation between model (1) and AS

incentive that looks like a return:

$$\frac{m^e}{p_1^e} = t_1^e + \frac{p_2^e}{p_1^e} t_2^e \quad \rightarrow \quad \bar{t}_1^e = t_1^e + \bar{R}_2^e t_2^e \quad (10)$$

The ratio $\bar{R}_2^e = p_2^e/p_1^e$ is the gross increase in the token price between the two periods. Unfortunately, this return applies to tokens, not to money as in (1).

Even though this is not the exact match we need, the normalized form (10) provides some useful clarifications that help to translate (9) into (1). First, it casts the seemingly atemporal money endowment m^e into a first-period resource: the normalized endowment \bar{t}_1^e . This is the maximum possible number of first-period tokens that could possibly be bought. It is analogous to the initial money endowment, but in tokens.

Second, the normalization shows that one of AS's decision variables is redundant. Performing the substitution $t_1^e = \bar{t}_1^e - \bar{R}_2^e t_2^e$ casts the problem purely in terms of future tokens. This is the variable that most closely resembles our saving decision. Helpfully, like the experimental saving constraint $0 \leq s_1^e \leq y_1^e$, the normalized equation ensures that the number of future tokens stays in bounds: $0 \leq t_2^e \leq \bar{t}_1^e/\bar{R}_2^e$.

Third, whenever it is necessary to operate in money units rather than tokens, the nor-

malization can be undone by multiplying by p_1^e throughout:

$$\begin{aligned} \bar{t}_1^e = t_1^e + \bar{R}_2^e t_2^e &\rightarrow p_1^e \bar{t}_1^e = p_1^e t_1^e + \bar{R}_2^e p_1^e t_2^e \\ 0 \leq t_2^e \leq \frac{\bar{t}_1^e}{\bar{R}_2^e} &\rightarrow 0 \leq p_1^e t_2^e \leq \frac{p_1^e \bar{t}_1^e}{\bar{R}_2^e} \end{aligned}$$

Once again, we have not made the obvious simplifications here. The reason is that this denormalization suggests an analogy to monetary saving in (1): $\bar{s}_1^e = p_1^e t_2^e$ is the current opportunity cost of buying t_2^e future tokens.

After making that substitution and simplifying, (9) becomes

$$\max_{\bar{s}_1^e} u\left((m^e - \bar{s}_1^e \bar{R}_2^e) - \omega_1^f\right) + \beta u\left(\bar{s}_1^e \bar{R}_2^e - \omega_2^f\right) \quad \text{s.t.} \quad 0 \leq \bar{s}_1^e \leq \frac{m^e}{\bar{R}_2^e}$$

Comparing this to (1), the exogenous experimental resources now match by setting $y_1^e = m^e$ and $y_2^e = 0$. The endogenous resources posed in \bar{R}_2^e terms, on the other hand, still do not have a natural analog. The money return in (1) appears in the future alone, but this token return appears in both periods.

The monetary implications of \bar{R}_2^e can be reconciled with (1) by normalizing $R_2^e = 1$ and setting $s_1^e = \bar{s}_1^e \bar{R}_2^e$:

$$\max_{s_1^e} u\left((m^e - s_1^e) - \omega_1^f\right) + \beta u\left(s_1^e - \omega_2^f\right) \quad \text{s.t.} \quad 0 \leq s_1^e \leq m^e \quad (11)$$

This formulation shows that AS, like AHLR, sets up a latent smoothing variable s_1^e . However, AS's version does not entail a money return ($r_2^e = 0$). The token return \bar{R}_2^e ultimately acts as an exchange rate between tokens and money, not as a return on saving. Importantly, in situations like this where $r_2^e = 0$, the participant's only reason to save is to smooth out the experimental windfall $y_1^e = m^e$. Saving has no investment use.

Formulation (11) underscores that AS's main source of experimental variation is m^e . Every m^e results in a unique saving amount s_1^e , no matter the token prices. (Given s_1^e , those

prices can be used to back out the token quantities.) Hence, changing the token prices while keeping m^e constant would simply pose the same smoothing question to the participant in different ways.

Because this task involves no risk, it activates only the intertemporal preference u and discounting β . It makes no statement about the other RU component, the risk preference ψ .

AS can be viewed as an extension of AHLR that allows participants to select their own saving amount, not just the extreme amounts $s_1^e \in \{0, y_1^e\}$. Interestingly, participants seem to prefer those extremes anyway: over half of Andreoni and Sprenger's decisions fall on the boundaries. Formulation (11) suggests two reasons this might occur.

First, participants cannot use experimental saving as an investment. This eliminates one of the main reasons to save, and can lead to $s_1^e = 0$. Second, if participants are on a skewed field consumption path, they will use experimental saving to force as many experimental resources as possible into the disadvantaged period. This can lead to either $s_1^e = 0$ or $s_1^e = y_1^e$, depending on which period needs more support. As we show in the numerical exercises below, the balance of field and experimental incentives needed to sustain an interior s_1^e in AS is actually quite delicate.

AS's field assets ω_1^f and ω_2^f , like AHLR's, are considered exogenous. The specification errors we discussed for AHLR therefore apply to AS as well, but the Stone-Geary form raises a new concern. Because Stone-Geary utility is not even defined before reaching the ω_1^f and ω_2^f consumption levels, those amounts are effectively exempted from smoothing. That is not the way consumption smoothing is usually understood. Indeed, when Andreoni and Sprenger estimate these quantities rather than imputing them, they do not always find negative values.

4 Robustness to Lifecycle Asset Integration

Having placed HL, AHLR, and AS into the framework of (1), we next examine their robustness to various integration assumptions. That analysis involves numerically changing

the levels of field resources, without paying much attention to what those resources mean. To provide some context in that regard, we begin by summarizing how the literature on background resources in the utility function has evolved to date.

4.1 Literature on Background Resources

In the theory of choice under static risk, decision makers evaluate their utility with respect to a terminal criterion. That framework conveniently compacts temporal problems into static ones whose goal resolves during a “final period.” In such settings, parameters like ω_1^f and ω_2^f are considered to be background wealth levels (Pratt, 1964; Arrow, 1971; Binswanger, 1981; Heinemann, 2008).

That formulation of utility carries some problematic implications. In critiques based on the terminal-wealth interpretation of EU, Hanssen (1988) and Rabin (2000) show that the assumption of reasonable risk aversion at low stakes implies absurd forms of risk aversion at high stakes. Both Rabin and Rabin and Thaler (2001) consider this inconsistency to be serious enough to warrant scrapping EU.

Cox and Sadiraj (2006) show that this troublesome issue does not arise if utility is evaluated with respect to *changes* in wealth. Those changes are usually interpreted as income. As a practical matter, incorporating those changes requires adding a wealth baseline to the utility argument.

Adopting that notion of background wealth, Andersen et al. (2018) investigate the relevance of asset integration in a static risk experiment. Their wealth baseline is defined quite broadly: it includes durable goods, real estate, and debt service; but not cash, equity in private companies, or non-tradeable assets. They find that participants integrate those baseline assets with experimental cash incentives quite weakly.

In the theory of lifecycle choices under risk, by contrast, ω_1^f and ω_2^f are often considered to be background consumption levels. Experiments with a temporal aspect usually assume that asset integration originates from background consumption (Cubitt and Read, 2007). For

example, AHLR and AS each take utility to be consumption-based rather than wealth-based, and they each include a background parameter in the utility argument.

Somewhat confusingly, both consumption- and wealth-based utility are amenable to time interpretations. For example, the multiperiod portfolio model has a terminal-wealth objective, while the multiperiod saving model has a lifecycle-consumption objective. Both have first-order conditions that take the canonical dynamic form $E_t(m_{t+1}R_{t+1}) \lesseqgtr 1$ discussed earlier. However, their respective outcomes are governed by very different utility features.

As a consequence, it is critical to be able to identify a model as consumption- or wealth-based. Eeckhoudt and Schlesinger provide a reliable way to do this, by observing what happens when an n^{th} -order risk is added to the utility argument. This will activate the n^{th} derivative if utility is wealth-based, and the $n + 1^{\text{th}}$ derivative if it is consumption-based. For EU ($u = \psi$), that reduces to determining whether the response to a second-order variance risk is governed by the Arrow-Pratt coefficient $-\psi''/\psi'$, or the Kimball coefficient $-\psi'''/\psi''$.

Meyer and Meyer (2005) discuss another tricky feature of the utility argument: its ability to create paradoxes in lifecycle models. For example, the equity-premium and riskfree-rate puzzles manifest as inconsistencies among intertemporal substitution, equity premia, and riskfree rates. These inconsistencies can be traced to artifacts of the utility specification.

The EU model at the root of the problem has isoelastic consumption utility and an isoelastic wealth value function. It also defines consumption quite broadly, while binding consumption tightly to wealth. That set of assumptions ultimately proves to be incompatible with US data.

When consumption is only a fraction of wealth (as in the US), consumption utility can be isoelastic only if wealth utility exhibits increasing elasticity, and vice versa. Keeping both aspects isoelastic requires defining wealth more narrowly than consumption. Because that redefinition would make wealth and consumption incoherent, Meyer and Meyer recommend using utility functions that are more flexible than isoelastic ones.

RU is one route to that flexibility. The version we use has a utility-of-wealth func-

tion that ranks risky consumption possibilities via the certainty equivalent, and a utility-of-consumption function that allocates the certainty equivalent across time. These two functions admit different risk and intertemporal elasticities by default. Those concepts could certainly be disentangled in other ways, but this method results in an intuitive SDF (3) where each attitude is governed by its own marginal rate of substitution.

4.2 Holt and Laury

HL implements ψ as the expo-power function⁸

$$\psi(c) = \frac{1}{\alpha_\psi} \left[1 - \exp\left(-\alpha_\psi \cdot \frac{c^{1-\rho_\psi}}{1-\rho_\psi}\right) \right]$$

Consistent with Meyer and Meyer’s stress on flexibility, the expo-power form permits increasing, decreasing, and constant RRA. Exponential utility ($\rho_\psi = 0$) and isoelastic power utility ($\alpha_\psi \rightarrow 0$) are special cases. Our baseline calibration uses Holt and Laury’s representative-agent estimates $\alpha_\psi = 0.03$ and $\rho_\psi = 0.73$, a combination that generates IRRA.

Exogenous Field Resources

A risk-neutral participant facing the baseline MPL in Table 1 would switch from the safe lottery to the risky one at line 5. Holt and Laury’s cohort switches further down the MPL on average, signalling risk aversion. But, the position of that switch point is highly sensitive to the the level of exogenous resources y_2^f .

Figure 1 shows how the baseline switch-point prediction changes as y_2^f increases. At HL’s assumed level $y_2^f = \$0$, the switch-point prediction is line 8.⁹ That prediction remains intact only up to $y_2^f = \$0.10$ of exogenous resources. It falls to line 7 by $y_2^f = \$0.20$, and to line 6 by $y_2^f = \$2$. Line 5 – full risk neutrality – occurs by $y_2^f = \$7$.

⁸HL’s expo-power function is actually slightly less flexible than this one (Xie, 2000). HL’s a and r parameters can be reconciled with α_ψ and ρ_ψ by setting $\alpha_\psi = a$ and $\rho_\psi = 1 - r$.

⁹HL also includes decision error by adding a discrete-choice framework on top of the structural comparison (5). To make our model’s predictions as crisp as possible, we abstract from these sorts of errors throughout.

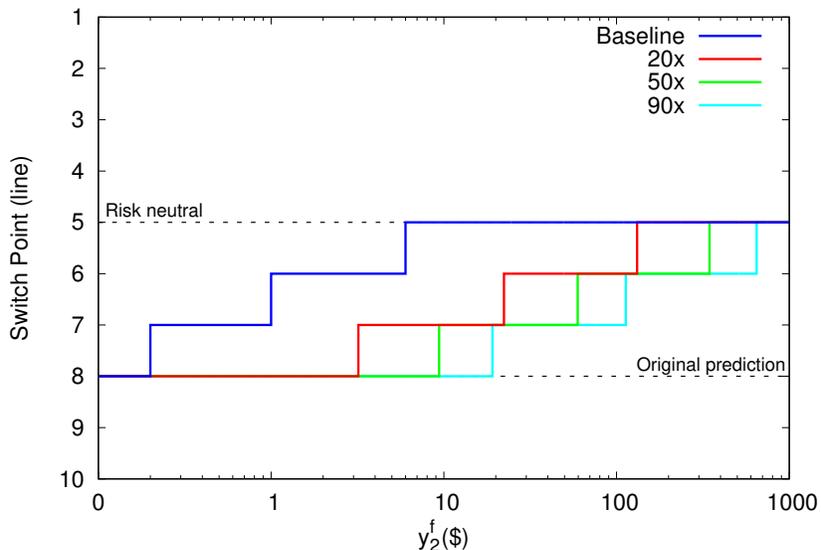


Figure 1: Predicted HL switch points as exogenous field resources increase

As a consequence, the original HL structural estimates are probably strongly predicated on the assumption of no asset integration. It is hard to imagine that participants did not consider the consequences of a mere \$0.20 of field resources during that task. If they did, they would have had to almost purposefully erect a mental divider between the experiment and the field.

HL also presents participants with a second MPL at a multiple of the original payoffs. Figure 1 repeats the baseline analysis at multiples of 20x, 50x, and 90x. This shows that the interference can be postponed, but not escaped, by scaling up the payoffs. Risk-neutral decisions eventually occur around $y_2^f = \$100$ at 20x, and around $y_2^f = \$700$ at 90x.

In exercises not shown here, we scale the exogenous field assets instead of the lottery stakes. We examine $y_2^f = \$1,000$ and $\$5,000$, two amounts that could easily reflect a household's monthly resources. Risk-neutral decisions occur all the way up to the 137x MPL in the first case, and up to the 611x MPL in the second.

HL's salience thus depends strongly on the relative levels of field and experimental resources. Because HL is a risk task, this finding can be partly contextualized within the Rabin critique. Namely, as exogenous field resources push the domain of experimental deci-

sion making to higher wealth levels, participants no longer make risky experimental decisions, even though they are truly risk averse. Unlike some of Rabin’s examples, however, the interference here does not require large or infinite amounts. The consequences manifest at low levels of field resources, levels that probably describe many real participants.

Endogenous Field Resources

To investigate the role of endogenous resources in HL, we must first parameterize the participant’s field smoothing environment. We consider the time interval to be monthly, so that field resources are relatively small. We set the baseline field return to $r_2^f = 1\%$ for a similar reason.

To keep s_1^f ’s behavior within the scope of existing theoretical results, we limit risks to mean-preserving spreads (MPS). We construct an MPS starting from a balanced income stream $y_1^f = E_1(\tilde{y}_2^f)$. We then create a two-outcome lottery centered around the future mean, which has 50-50 probabilities. The lottery payments are thus $E_1(\tilde{y}_2^f) \pm \Delta y_2^f$, where Δy_2^f is the spread.

We must also flesh out the other two preference domains. We parameterize the discounting parameter with $\beta = 0.999$, which corresponds to an annual parameter of 0.988. We implement the smoothing preference u with another expo-power function

$$u(c) = \frac{1}{\alpha_u} \left[1 - \exp\left(-\alpha_u \cdot \frac{c^{1-\rho_u}}{1-\rho_u}\right) \right]$$

The analog to RRA for u is the relative resistance to intertemporal substitution (RRIS), or inverse EIS. We set the baseline parameters to $\alpha_u = 0$ and $\rho_\psi = 2$, a power utility function in the neighborhood of macroeconomic estimates.

We illustrate the interference from endogenous field resources using the 20x MPL. We set the field income levels to $y_1^f = E_1(\tilde{y}_2^f) = \100 , and omit field risk for clarity ($\Delta y_2^f = 0$). We assume first that no field smoothing instrument exists, and then that a field smoothing

Line	No field saving		$r_2^f = 1\%$		$r_2^f = 10\%$	
	$s_1^{f, safe}$	$s_1^{f, risky}$	$s_1^{f, safe}$	$s_1^{f, risky}$	$s_1^{f, safe}$	$s_1^{f, risky}$
1	0	0	33.60	46.66	33.50	45.71
2	0	0	33.22	44.32	33.15	43.46
3	0	0	32.84	41.81	32.79	41.07
4	0	0	32.45	39.10	32.43	38.49
5	0	0	32.06	36.16	32.06	35.69
6	0	0	31.66	32.91	31.68	32.61
7	0	0	31.26	29.25	31.30	29.15
8	0	0	30.84	24.99	30.92	25.13
9	0	0	30.43	19.68	30.53	20.18
10	0	0	30.00	11.64	30.13	12.91
Switch Point	8		6		6	

Table 6: HL 20x decisions with and without a field saving instrument ($y_1^f = y_2^f = \$100$)

instrument exists and pays returns of $r_2^f = 1\%$ and 10% .

Table 6 presents the switch points under each assumption, as well as the hidden field saving amounts $s_1^{f, safe}$ and $s_1^{f, risky}$ that arise on each line of the MPL. Several features of this table are important. First and foremost, just by assuming that the participant has some form of field smoothing, the switch-point prediction rises from line 8 to line 6. This occurs because \tilde{y}_2^e represents a second-period windfall that the participant would like to smooth back to the first period, but that option was previously unavailable.

To reiterate, when field smoothing is impossible, the participant tolerates the second period’s lopsided but risk-averse consumption outcome. But, when field smoothing is allowed, the participant seemingly becomes *less* risk averse. This is very strange: we have not touched ψ (or any other preference domain), and so we know for a fact that the participant’s risk attitude is the same in both cases.

This movement up the MPL occurs because field smoothing allows the participant to self-insure by moving funds from the first period to the second. This “hedge across time” ends up being much more powerful than the “hedge within time” afforded by static risk aversion. The ability to self-insure with lifecycle resources ultimately allows the participant to engage in more MPL risk.

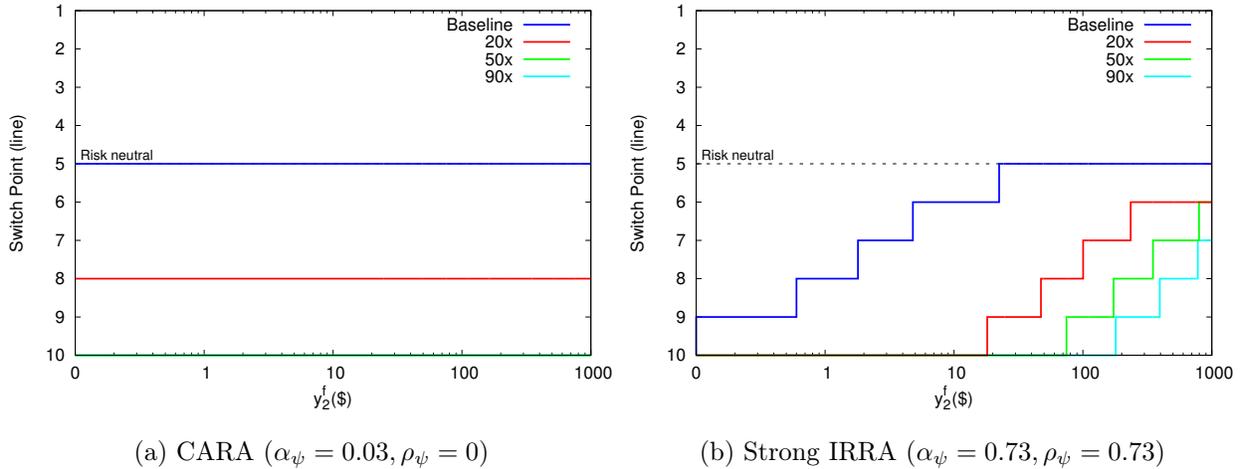


Figure 2: Predicted HL switch points for other ψ specifications

Second, as our theoretical discussion suggested, each side of the utility comparison (6) does indeed entail a different amount of field saving. The increase from $s_1^{f, safe}$ to $s_1^{f, risky}$ on each line is due to the classic smoothing and precautionary responses. Specifically, the change in the smoothing response is triggered by the difference in $\tilde{y}_2^{e, safe}$'s and $\tilde{y}_2^{e, risky}$'s means, and the change in the precautionary response is triggered by the difference in their variances.

Third, for each value of r_2^f , s_1^f always decreases moving down the MPL. This occurs because the highest payoff becomes increasingly more assured, and so less field saving is needed to guarantee a good consumption path.

Thus, HL decisions are also sensitive to whether the participant has a field smoothing instrument. A participant who seems to be risk neutral per the MPL may be truly risk averse, but self-insuring outside the experimenter's view.

Other Considerations

Figure 2 re-examines interference from exogenous field resources under two rather extreme risk attitudes. These exercises illustrate how α_ψ and ρ_ψ act together to explain decisions at very different payoff scales.

The left panel plots decisions under CARA. For this attitude, the only aspect of risk

that matters is the level of \tilde{y}_2^e 's payoffs. Because that level is fixed for a given scaling, and because y_2^f is not risky, the switch point within a scaling never varies with y_2^f .

Under this CARA parameterization, the switch point at the baseline scaling is line 5, completely indistinguishable from risk neutrality. The 20x scaling now yields the switch point at line 8. And, no switches ever occur at 50x and 90x. CARA thus implies a huge amount of risk aversion at large stakes.

The right panel repeats this exercise under much stronger IRRA. Unlike the CARA results, these switch points do eventually change as y_2^f increases. Thus, ρ_ψ 's role in the expo-power function is to tamp down α_ψ 's explosive tendencies at large amounts, and to provide some risk aversion at small amounts.

4.3 Andersen et al.

Andersen et al.'s baseline estimates entail an isoelastic u with $\rho_u = 0.74$, and an annual pure rate of time preference of 10%.

We illustrate the interference from lifecycle asset integration using the annual percentage rate (APR) structure in Table 3 at intervals of 1, 3, 6, and 12 months. We construct a two-period environment by breaking those intervals into two equally long segments.

To keep the units uniform across our examples, we convert DKK to USD at 6.55 to 1, Andersen et al.'s reported exchange rate. This places y_1^e at about \$450. Because u is isoelastic, multiplicatively scaling its argument by the exchange rate does not affect any of the switch-point predictions.

Exogenous Field Resources

Figure 3 presents the switch-point predictions as exogenous field resources increase. Andersen et al. calibrate their exogenous field resources to a government consumption survey that finds average daily field consumption to be DKK 118. In Figure 3's units, one day's worth of background consumption is about \$20, and three months' worth about \$1,600.

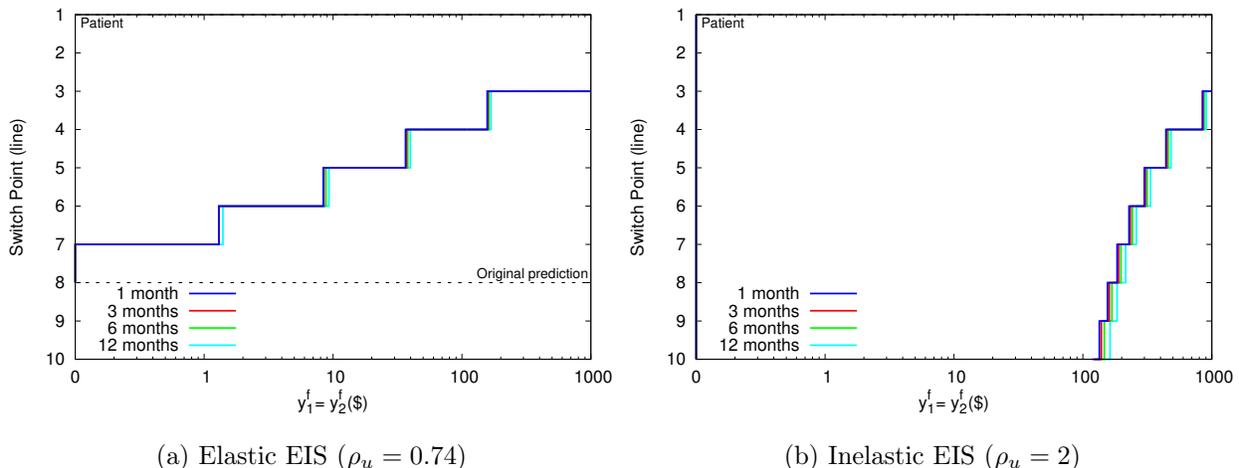


Figure 3: Predicted AHLR switch points as exogenous field resources increase

The switch points in the left panel are certainly sensitive to y_1^f and y_2^f . The switch point moves from line 8 at $y_1^f = y_2^f = \$0$ to line 3 at $y_1^f = y_2^f = \$1,000$. That prediction evolves in essentially the same way at different τ intervals.

Because u is estimated in conjunction with a supplemental HL task, it is not entirely surprising that Andersen et al. find ρ_u to be close to Holt and Laury’s ρ_ψ (0.74 vs. 0.73). But, it is important to remember that those parameters address different attitudes: ρ_u measures u ’s RRIS, while ρ_ψ measures ψ ’s RRA. Thus, the same numerical estimate signals something different in AHLR’s intertemporal context than it does in HL’s static context. For AHLR, this estimate means that RRIS is 0.74, and so EIS is $1/0.74 = 1.35$.¹⁰

Problematically, elastic EIS is hardly ever found in macroeconomic data. So, the right panel of Figure 3 repeats this exercise under $\rho_u = 2$, the same inelastic value used in the HL analysis. These switch-point predictions move much more sharply than the elastic ones. They do not even dislodge from the bottom of the MPL until about $y_1^f = y_2^f = \$100$, but they still reach line 3 by $y_1^f = y_2^f = \$1,000$. There is slightly more separation in the predictions at different time intervals, but they largely move in tandem as before.

¹⁰This distinction between RRA and RRIS cannot be extinguished by forcing them to be “mathematically equivalent” by setting $u = \psi$. In this intertemporal context, that equivalence simply collapses RU to EU. It does not somehow compress RRA and RRIS into a single attitude. Both of those elasticities are still present under EU – they just take the same value.

Line	No field saving		$r_2^f = 1\%$		$r_2^f = 10\%$	
	$s_1^{f,early}$	$s_1^{f,late}$	$s_1^{f,early}$	$s_1^{f,late}$	$s_1^{f,early}$	$s_1^{f,late}$
1	0	0	209.37	-252.19	240.82	-201.37
2	0	0	209.37	-258.14	240.82	-206.76
3	0	0	209.37	-264.16	240.82	-212.21
4	0	0	209.37	-270.26	240.82	-217.72
5	0	0	209.37	-276.43	240.82	-223.31
6	0	0	209.37	-282.67	240.82	-228.96
7	0	0	209.37	-288.99	240.82	-234.67
8	0	0	209.37	-295.38	240.82	-240.45
9	0	0	209.37	-301.84	240.82	-246.30
10	0	0	209.37	-308.37	240.82	-252.21
Switch Point	3		1		4	

Table 7: AHLR 6-month decisions with and without a field saving instrument ($y_1^f = y_2^f = \$500$)

Importantly, many switch-point predictions can be rationalized under either elasticity assumption. For example, line 5 arises at about $y_1^f = y_2^f = \$10$ under the elastic u , and at about $y_1^f = y_2^f = \$300$ under the inelastic u . This illustrates that AHLR, like HL, is sensitive to the level of exogenous resources. In this case, as field resources increase, the participant appears to be more patient.

Unlike HL, this loss of salience cannot be attributed to the Rabin critique. Because AHLR has no risk, it evades that concern entirely. Here, the loss is wholly a consequence of the participant using the experiment to smooth out the field consumption path. When field resources are large enough, the desired smoothness always results in choosing the early option.

Endogenous Field Resources

Table 7 presents the switch-point predictions for the six-month MPL, with and without the assumption that the participant has a field smoothing instrument. We again set field returns to $r_2^f = 1\%$ and 10% . We set field resources to $y_1^f = E_1(y_2^f) = \$500$, and continue to omit field risk ($\Delta y_2^f = 0$).

As with HL, the AHLR switch-point predictions change just by assuming that the participant has a field smoothing instrument. The no-field-smoothing prediction is line 3. This rises to line 1 when $r_2^f = 1\%$, and actually falls to line 4 when $r_2^f = 10\%$.

Also like HL, the amount of field smoothing changes markedly on each side of the AHLR comparison (8). The value of s_1^f is always positive for the early option, indicating that the participant will smooth forward some of the early experimental windfall. Similarly, s_1^f is always negative for the late option, indicating that the participant will smooth back some of the late windfall.

The move from line 3 to line 1 when $r_2^f = 1\%$ makes the participant appear more patient. But, because we have not changed anything about preferences, we know that this outcome arises from some sort of substitution. In this case, the ability to smooth in the field allows the participant to support a seemingly more patient outcome in the MPL. That “patience” is nothing more than opting for the larger late experimental payment, and smoothing some of it back with the field instrument.

Those same smoothing attitudes are present when $r_2^f = 10\%$. But, this situation raises another consideration: the high return on field saving makes it a very attractive investment. That countervailing factor causes the participant to opt for the earlier experimental payment, and save it with the field instrument. The switch point therefore moves down the MPL instead of up.

So, the mere ability to smooth in the field can also result in different AHLR predictions. Saving’s smoothing use can be triggered with fairly small field returns. Larger returns can also pull in saving’s investment use.

Other Considerations

Even though AHLR involves no experimental risk, a participant could still be exposed to field risk. We explore those implications by again setting $y_1^f = E_1(y_2^f) = \$500$, and varying the MPS spread Δy_2^f from \$0 to \$500. For clarity, we assume away endogenous field resources.

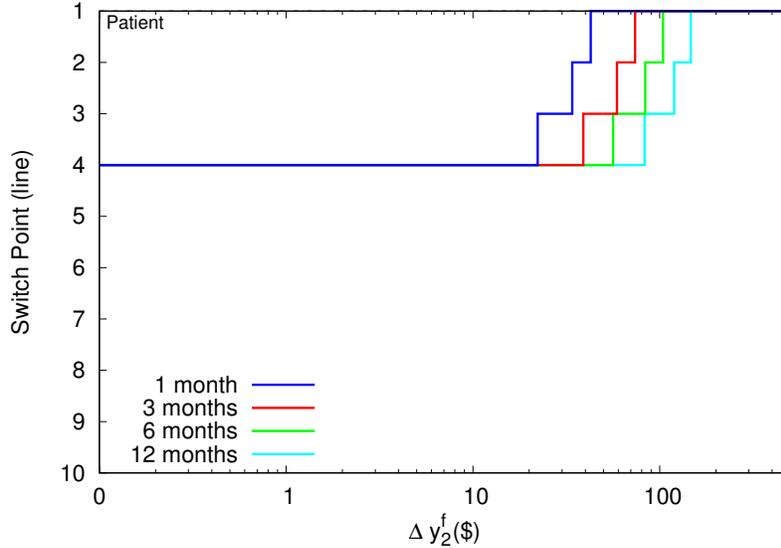


Figure 4: Predicted AHLR switch points under RU, for MPS of \tilde{y}_2^f ($y_1^f = E_1(\tilde{y}_2^f) = \500)

Because this set of incentives will activate both risk preferences and smoothing preferences, we use a full RU specification. We set ψ to Holt and Laury’s IRRA estimates, and u to the previous power function with $\rho_u = 2$. To reiterate our earlier note on the difference between RRA and RRIS, it would be unwise to impute u with Holt and Laury’s estimates. Those values would imply IRRIS, which would eventually send the EIS all the way to 0.

Figure 4 plots the resulting switch points. As Δy_2^f increases, the switch point moves up the MPL. Because this exercise involves pure changes in risk, that movement comes from the participant’s precautionary motive alone. The precautionary motive, as a rule, offsets higher future risk with more saving. In the context of AHLR, “more saving” means “choosing the late option.”

So, because participants can potentially use experimental smoothing to mitigate field risk, the experimenter should also have a good grasp on how much field risk participants face. Higher field risk will lead to more experimental smoothing.

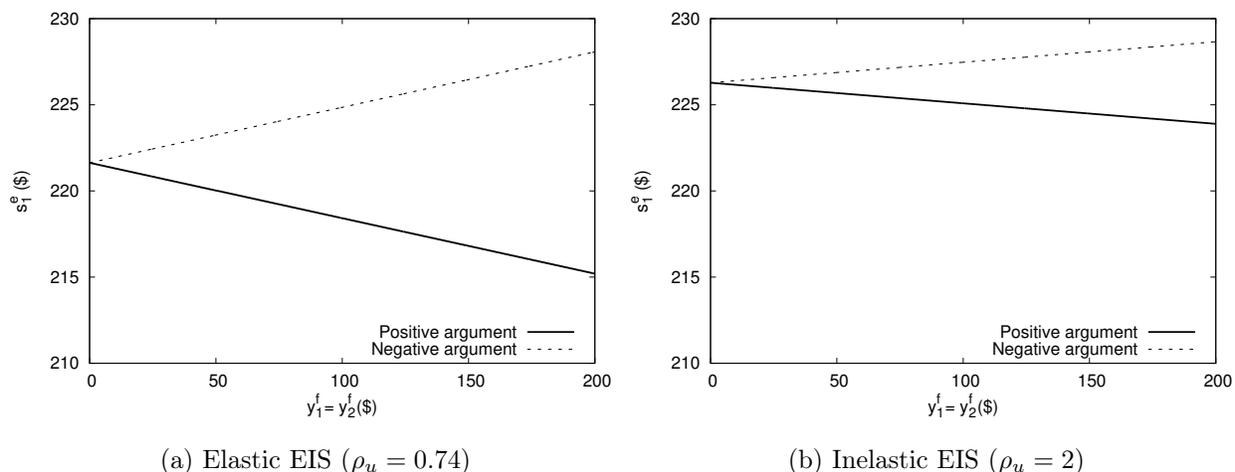


Figure 5: Predicted AS saving amounts when exogenous field resources are construed as positive or negative

4.4 Andreoni and Sprenger

Because the nature of AS interference does not depend on Andreoni and Sprenger’s specific m^e values, we keep AHLR’s payoff structure for the AS exercises. In other words, we ask what decisions AHLR participants would have made if the incentives had been presented as an AS task with $m^e = \text{DKK } 3,000$. This allows us to easily compare the AS findings to the results above.

Exogenous Field Resources

An obvious difference between AHLR and AS is AS’s Stone-Geary treatment of exogenous field resources. Figure 5 plots the predicted AS saving amounts s_1^e assuming that y_1^f and y_2^f can be positive or negative. This shows that Stone-Geary inverts the standard relationship between saving and field resources.

The downward-sloping lines reflect the usual understanding of that relationship, where the propensity to save falls as the participant holds additional positive resources. This occurs because those extra resources push the consumption interval away from the highly curved parts of u , thereby diminishing the desire to smooth.

The upward-sloping lines reflect the Stone-Geary conception. Because the movements along the utility function now operate in reverse, the propensity to save rises. As the participant receives additional negative resources, a greater part of the consumption interval falls into the very curved areas of u . Because those areas reflect undesirably lumpy outcomes, the participant saves more to smooth out the consumption path.

Because AS entails $r_2^e = 0$, experimental saving has no investment use here. Figure 5 shows that participants will save anyway. This illustrates that saving does indeed have a meaningful smoothing function separate from its investment function. Unfortunately, Figure 5 also shows that the field resources y_1^f and y_2^f can once again interfere with the experimental decision s_1^e . As the participant's field resources increase, the desire to save in the experiment falls.

Endogenous Field Resources

Figure 6 repeats this exercise assuming that the participant has field smoothing instruments that return $r_2^f = 1\%$, 5% , and 10% . These plots provide the clearest illustration yet of the elasticity (4) between experimental and field smoothing: whenever s_1^e falls, s_1^f simultaneously rises to compensate.

Because r_2^f is always greater than $r_2^e = 0$ in Figure 6, that tradeoff always moves in the direction of substituting experimental saving with field saving. Importantly, the fact that field saving has a higher return does not necessarily mean that the participant will forego experimental saving. Instead, up to about $y_1^f = y_2^f = \$200$, the participant saves using both instruments. That mixture allows the participant to trade some experimental saving for better-returning field saving, while also keeping the consumption path sufficiently smooth.

Past $y_1^f = y_2^f = \$200$, however, the participant does forego experimental saving. In fact, the participant would like to borrow from the experimenter and save that money in the field, but the rules forbid doing so. As a consequence, a participant with more than \$200 of field resources will always choose $s_1^e = 0$ in AS. Importantly, the experimenter will

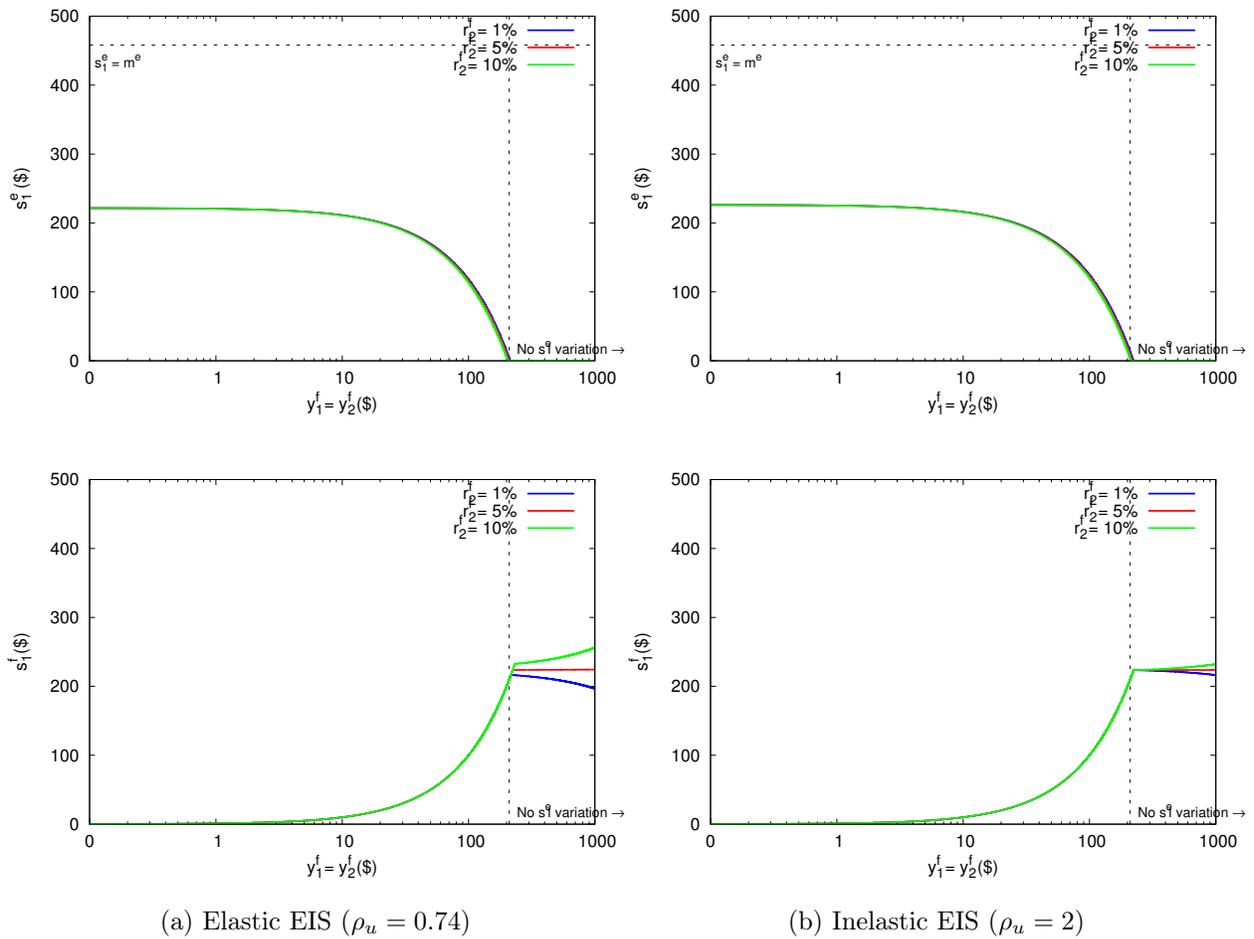
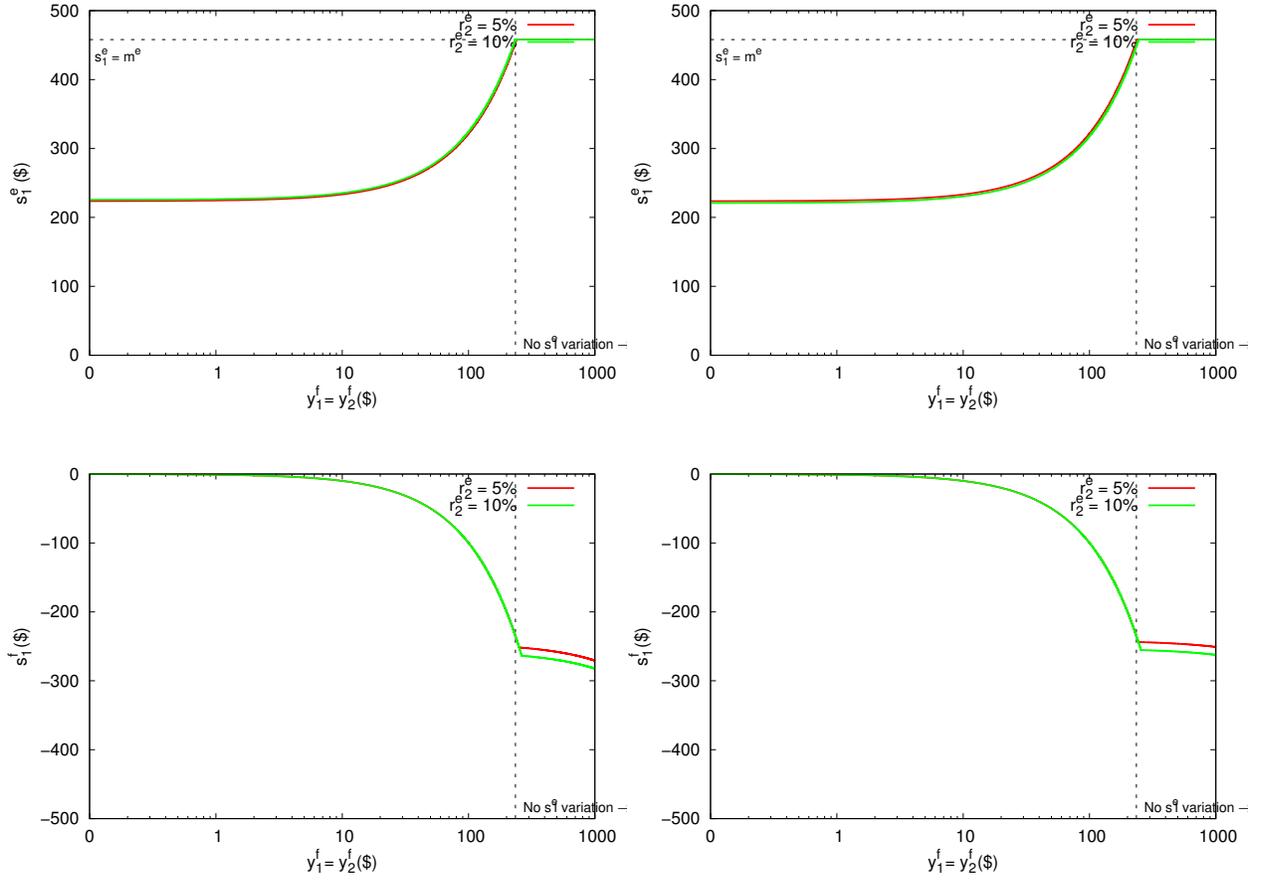


Figure 6: Predicted AS saving amounts when a field saving instrument is present



(a) Elastic EIS ($\rho_u = 0.74$)

(b) Inelastic EIS ($\rho_u = 2$)

Figure 7: Predicted AS saving amounts if s_1^e has a return ($r_2^f = 1\%$)

not be able to determine *ex post* whether that boundary outcome has been generated by the participant's preferences (e.g., strong impatience or high RRIS) or by this interference. Those explanations are observationally equivalent.

Elasticity (4) can cause problems at the upper bound just as easily. Figure 7 contains a similar exercise assuming that the AS task could be modified to pay experimental returns $r_2^e = 5\%$ and 10% , while the participant's field return is $r_2^f = 1\%$. This change makes the experiment a better investment than the field, and so the MRS operates in reverse. Experimental saving now reaches its upper bound $s_1^e = m^e$ at about $y_1^f = y_2^f = \$200$. The participant compensates for higher experimental saving by borrowing in the field.

In sum, AS is sensitive to endogenous field resources in a way that naturally results

in boundary decisions. Notably, interference from these resources can result in higher or lower experimental saving. The former is like the attenuation towards patient outcomes that occurs with AHLR.

5 Discussion

This paper overlays a thought experiment onto the design of risk and time experiments: what would happen if participants treated an experiment not as a wholly isolated task, but as an injection of resources into their existing lifecycle plans? In the three classic designs we examine – encompassing static and dynamic tasks, risk and time manipulations, as well as menu and free responses – experimental decisions change markedly when participants are permitted to trade off experiment and field resources, versus not. In other words, inferences in these experiments are not robust to lifecycle asset integration. Moreover, because field interference occurs reliably, and with small amounts, we suspect the inference problem extends beyond these three settings.

We reach those conclusions by wrapping each experiment into a meta-model that represents a participant’s unified experiment-field environment. The underlying two-period framework is a workhorse in intertemporal choice theory, with well-studied properties. It allows us to sidestep a major confound that would arise in any empirical examination of lifecycle asset integration: participants’ preferences. Within our theoretical sandbox, we can absolutely ensure that preferences are held constant while field resources change, and thereby establish definitively that lifecycle asset integration causes interference.

In reality, experimenters will not be able to observe the mental act of asset integration. They will therefore be unable to detect interference from just a cursory examination of their data. Helpfully, our theoretical framework provides a way to envision how this unobservable might impact many tasks.

The inference problem can be seen with a simple AHLR illustration. Suppose that an

experimenter observes switchpoints at rows 5, 6, and 7 of the AHLR task menu. Under the standard interpretation of those data, the row-5 switchpoint reflects the most patient preferences, the row-7 switchpoint reflects the least patient preferences, and the row-6 switchpoint reflects preferences in between. In other words, switching further down the menu implies less patience.

Once lifecycle asset integration is admitted as a possibility, however, the experimenter faces a quandary. Per our analysis, those same switchpoints can now be interpreted in three conflicting ways:

1. If participants have not integrated anything, the standard interpretation of rows 5-7 is correct.
2. If participants have integrated exogenous field resources, rows 5-7 will reflect preferences that are less patient than the standard interpretation. Very problematically, the straightforward relationship between rows and patience will almost surely break if different participants have integrated different amounts. For example, a participant with low patience and moderate integration could switch at row 5, while one with similarly low patience but little integration switches at row 7 as before. Participants' intertemporal preferences, which smooth field and experimental resources collectively, are responsible for this divergence.
3. If participants have integrated endogenous field resources, rows 5-7 will reflect preferences that are either more or less patient than the standard interpretation, depending on the size of the field return. When that return is small, issues like those stemming from exogenous resources will occur. When that return is large, rows further down the table might actually imply more patience. For example, a participant with high patience and moderate integration could switch at row 7, while one with similarly high patience but little integration switches at row 5 as before. Participants' intertemporal preferences are responsible once again, but the intertemporal tradeoffs are complicated

by the fact that the field return makes a very good investment.

These interpretations are all observationally equivalent: experimenters cannot tell whether the switchpoint data have been generated by the preferences targeted by the task, or by lifecycle asset integration. We find similar instances of observational equivalence in the static HL risk task, and in the free-response AS saving task. Field resources on the scale of \$100 to \$1,000 almost always cause problems, but resources on the scale of \$10 frequently do too. The usual outcome is a kind of attenuation bias that pulls data toward risk neutral and patient interpretations, even though preferences are solidly risk averse and impatient.

The lone bit of *prima facie* evidence for lifecycle asset integration appears to be large numbers of boundary decisions. Because asset integration involves substituting between experimental and field resources, experimental incentives that are quite stingy or quite rich relative to the field can result in minimal or maximal uptake. It may be tricky to elicit interior decisions in free-response designs as a consequence. In the past, we ourselves have dismissed large numbers of boundary decisions in experimental runs as some sort of generic design flaw. In light of these results, that flaw is likely very specific.

It is also not fatal, at least conceptually. If experimenters suspect that asset integration will occur, they can implement mitigation strategies both *ex ante* during design, and *ex post* during statistical analysis. *Ex ante*, experimenters can pre-screen participants to appraise field resources. Doing so gives experimenters a better chance of matching their incentives to participants' lifecycle assets. Incentives that match poorly are predestined to yield little data variation, due to lifecycle asset integration's attenuating effects.

Ex post, experimenters should bear in mind that this attenuation reduces the sensitivity of statistical tests. It will thus shrink *a priori* the differentials underpinning basic comparison tests, and promote false negatives. To reiterate, the data generating those differentials cannot by themselves reveal whether the attenuation is enough to ruin a test.

For studies that estimate preference parameters, the structural equations should start from something like the joint system (2a) and (2b), which captures the entire experiment-field

interaction.¹¹ This is a much taller order than the oft-used method of eliciting preferences on a “one task per attitude” basis, and then uniting those tasks into a structural likelihood resembling

$$\text{likelihood} = \text{smoothing likelihood} + \text{risk-aversion likelihood} + \text{discounting likelihood}$$

Besides omitting the field, this method weighs all three preferences equally, but those preferences do not actually have equal weights in the decision model (1) generating the data. SDF (3), embedded in (2a) and (2b), inherently provides the correct weighting.¹²

Many structural models assume EU. In our experience, this opens up the experimenter to confusing RRA and RRIS, particularly when an assortment of tasks on risk and time has been assembled together. Helpfully, the RU preferences in (1) make plain that some descriptions of risk aversion just do not make sense as elasticities of intertemporal substitution. For example, Holt and Laury’s IRRRA estimate would not make a good RRIS, because IRRIS eventually implies no intertemporal elasticity.

The benefits of an RU stance toward modeling are not rooted in the RU functional form *per se*, but in contextualizing experimental design within the trio of discounting, risk aversion, and intertemporal substitution by default. Lifecycle asset integration will activate all of those preferences at once – no matter what preference the task is ostensibly designed to trigger – and RU-style framing faithfully captures that activation pattern. RU-style framing also clarifies which field data would serve as appropriate controls for the resulting experimental interference. If, after viewing an experiment in this light, the experimenter determines that simplifications are warranted (e.g., EU or a smaller set of field resources), those can certainly be made.

Each of our examples is a “laboratory experiment” in the sense of its having been orig-

¹¹When an experiment does not include anything resembling a smoothing decision, a single structural equation like (2a) is enough.

¹²In a numerical decomposition of saving into its smoothing motive (controlled by u) and precautionary motive (controlled by ψ), Bostian and Heinzl (2018) find that, under a wide range of risky scenarios, more than 80% of saving is attributable to smoothing.

inally conducted in a laboratory. Because laboratory experiments cleanly segregate the experimenter’s manipulation from all other incentives, those examples fit neatly into our meta-model. At the same time, that model’s explicit link between experiment and field calls out an uncomfortable limitation of the laboratory paradigm: a laboratory room does not necessarily wall off the field. Unquestionably, participants can mentally incorporate their own field resources into a task conducted inside a laboratory.¹³

This, in turn, suggests that the standard taxonomy’s distinctions between laboratory, field, artefactual, and natural experiments (Harrison and List, 2004) start to blur for risk and time designs. Indeed, our analysis shows that trying to make those designs “more realistic” with large, field-scale incentives does not change anything. Field-scale lifecycle asset integration will interfere in those situations too. Because interference can occur at any scale, experimenters should have a good understanding of their participants’ field resources when implementing these tasks. Neglecting the diversity of lifecycle assets across cohorts could be a source of replication failure (Camerer et al., 2016).

The notable prior successes of these experiments may be fortuitously linked to their subject pools, which usually contain students or other groups with small amounts of lifecycle assets. Our model contains this scenario as a special case, and grants that designs will largely work as planned when field resources are low. But, the model also stresses that the external validity of this special case should not be reflexively inflated, because lifecycle asset integration can intrude easily.

References

- Andersen, Steffen, Harrison, Glenn W., Lau, Morten I., and Rutström, E. Elisabet. Eliciting risk and time preferences. *Econometrica*, 76(3): 583–618, 2008.
- Andersen, Steffen, Harrison, Glenn W., Lau, Morten I., and Rutström, E. Elisabet. Discounting behavior: a reconsideration. *European Economic Review*, 71: 15–33, 2014.

¹³These examples also use financial incentives exclusively, but the relevant incentives might not always be financial. For instance, a farmer participating in an experiment in a developing country might use grain storage as a field smoothing instrument, instead of money.

- 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. *Review of Economics and Statistics*, 100(5): 816–830, 2018.
- Andreoni, James and Sprenger, Charles. Estimating time preferences from convex budgets. *American Economic Review*, 102(7): 3333–3356, 2012.
- Arrow, Kenneth J. *Essays in the Economics of Risk-Bearing*. Markham Publishing, Chicago, IL, 1971.
- Ballinger, T. Parker, Palumbo, Michael G., and Wilcox, Nathaniel T. Precautionary saving and social learning across generations: an experiment. *Economic Journal*, 113(490): 920–947, 2003.
- Ballinger, T. Parker, Hudson, Eric, Karkoviata, Leonie, and Wilcox, Nathaniel T. Saving behavior and cognitive abilities. *Experimental Economics*, 14(3): 349–374, 2011.
- Bansal, Ravi and Yaron, Amir. Risks for the long run: a potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4): 1481–1509, 2004.
- Binswanger, Hans P. Attitudes toward risk: theoretical implications of an experiment in rural India. *Economic Journal*, 91(364): 867–890, 1981.
- Bostian, AJ A. and Heinzl, Christoph. Precautionary saving under recursive preferences. Working paper, 2018.
- Brown, Alexander L., Chua, Zhikang Eric, and Camerer, Colin F. Learning and visceral temptation in dynamic saving experiments. *Quarterly Journal of Economics*, 124(1): 197–231, 2009.
- Camerer, Colin F., Dreber, Anna, Forsell, Eskil, Ho, Teck-Hua, Huber, Jürgen, Johannesson, Magnus, Kirchler, Michael, Almenberg, Johan, Altmejd, Adam, Chan, Taizan, Heikensten, Emma, Holzmeister, Felix, Imai, Taisuke, Isaksson, Siri, Nave, Gideon, Pfeiffer, Thomas, Razen, Michael, and Wu, Hang. Evaluating replicability of laboratory experiments in economics. *Science*, 351(6280): 1433–1436, 2016.
- Carroll, Christopher D. and Kimball, Miles S. Liquidity constraints and precautionary saving. NBER Working Paper 8496, 2005.
- Coble, Keith H. and Lusk, Jayson L. At the nexus of risk and time preferences: an experimental investigation. *Journal of Risk and Uncertainty*, 41(1): 67–79, 2010.
- Cochrane, John H. *Asset Pricing*. Princeton University Press, revised edition, 2005.
- Coller, Maribeth and Williams, Melonie B. Eliciting individual discount rates. *Experimental Economics*, 2(2): 107–127, 1999.
- Cox, James C. and Sadiraj, Vjollca. Small- and large-stakes risk aversion: Implications of concavity calibration for decision theory. *Games and Economic Behavior*, 56(1): 45–60, 2006.

- Cubitt, Robin P. and Read, Daniel. Can intertemporal choice experiments elicit time preferences for consumption? *Experimental Economics*, 10(4): 369–389, 2007.
- Drèze, Jacques H. and Modigliani, Franco. Épargne et consommation en avenir aléatoire. *Cahiers du Séminaire d'Économétrie*, 9: 7–33, 1966.
- Drèze, Jacques H. and Modigliani, Franco. Consumption decisions under uncertainty. *Journal of Economic Theory*, 5(3): 308–335, 1972.
- Eeckhoudt, Louis and Schlesinger, Harris. Changes in risk and the demand for saving. *Journal of Monetary Economics*, 55(7): 1329–1336, 2008.
- Epper, Thomas, Fehr, Ernst, Fehr-Duda, Helga, Kreiner, Claus Thustrup, Lassen, David Dreyer, Leth-Petersen, Søren, and Rasmussen, Gregers Nytoft. Time discounting and wealth inequality. *American Economic Review*, 110(4): 1177–1205, 2020.
- Epstein, Larry G. and Zin, Stanley E. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica*, 57(4): 937–969, 1989.
- Epstein, Larry G. and Zin, Stanley E. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: an empirical analysis. *Journal of Political Economy*, 99(2): 263–286, 1991.
- Filiz-Ozbay, Emel, Guryan, Jonathan, Hyndman, Kyle, Kearney, Melissa, and Ozbay, Erkut Y. Do lottery payments induce savings behavior? Evidence from the lab. *Journal of Public Economics*, 126: 1–24, 2015.
- Gollier, Christian. *The Economics of Risk and Time*. MIT Press, Cambridge, MA, 2001.
- Hanssen, Bengt. Risk aversion as a problem of conjoint measurement. In Gärdenfors, Peter and Sahlin, Nils-Eric, editors, *Decision, Probability, and Utility*, pages 136–158. Cambridge University Press, New York, NY, 1988.
- Harrison, Glenn W. and List, John A. Field experiments. *Journal of Economic Literature*, 42(4): 1009–1055, 2004.
- Harrison, Glenn W., Lau, Morten I., Ross, Don, and Swarthout, J. Todd. Small stakes risk aversion in the laboratory: a reconsideration. *Economics Letters*, 160: 24–28, 2017.
- Heinemann, Frank. Measuring risk aversion and the wealth effect. In Harrison, Glenn W. and Cox, James C., editors, *Research in Experimental Economics*, volume 12, pages 293–313. Emerald, Bingley, UK, 2008.
- Holt, Charles A. and Laury, Susan K. Risk aversion and incentive effects. *American Economic Review*, 92(5): 1644–1655, 2002.
- Ingersoll, Jonathan E. *Theory of Financial Decision Making*. Rowman&Littlefield Publishers, Savage, MD, 1987.

- Kahneman, Daniel and Tversky, Amos. Prospect theory: an analysis of decision under risk. *Econometrica*, 47(2): 263–292, 1979.
- Kimball, Miles and Weil, Philippe. Precautionary saving and consumption smoothing across time and possibilities. *Journal of Money, Credit and Banking*, 41(2-3): 245–284, 2009.
- Kimball, Miles S. Precautionary saving in the small and in the large. *Econometrica*, 58(1): 53–73, 1990.
- Kreps, David M. and Porteus, Evan L. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1): 185–200, 1978.
- Leland, Hayne E. Saving and uncertainty: the precautionary demand for saving. *Quarterly Journal of Economics*, 82(3): 465–473, 1968.
- Meyer, Donald J. and Meyer, Jack. Risk preferences in multi-period consumption models, the equity premium puzzle, and habit formation utility. *Journal of Monetary Economics*, 52(8): 1497–1515, 2005.
- Miao, Bin and Zhong, Songfa. Risk preferences are not time preferences: separating risk and time preference: comment. *American Economic Review*, 105(7): 2272–2286, 2015.
- Pratt, John W. Risk aversion in the small and in the large. *Econometrica*, 32(1–2): 122–136, 1964.
- Rabin, Matthew. Risk aversion and expected utility theory: a calibration theorem. *Econometrica*, 68(5): 1281–1292, 2000.
- Rabin, Matthew and Thaler, Richard H. Anomalies: risk aversion. *Journal of Economic Perspectives*, 15(1): 219–232, 2001.
- Rabin, Matthew and Weizsäcker, Georg. Narrow bracketing and dominated choices. *American Economic Review*, 99(4): 1508–1543, 2009.
- Read, Daniel, Lowenstein, George, and Rabin, Matthew. Choice bracketing. *Journal of Risk and Uncertainty*, 19(1–3): 171–197, 1999.
- Rothschild, Michael and Stiglitz, Joseph E. Increasing risk II: its economic consequences. *Journal of Economic Theory*, 3(1): 66–84, 1971.
- Sandmo, Agnar. The effect of uncertainty on saving decisions. *Review of Economic Studies*, 30(3): 353–360, 1970.
- Schechter, Laura. Risk aversion and expected-utility theory: a calibration exercise. *Journal of Risk and Uncertainty*, 35(1): 67–76, 2007.
- Selden, Larry. A new representation of preferences over “certain x uncertain” consumption pairs: the “ordinal certainty equivalent” hypothesis. *Econometrica*, 46(5): 1045–1060, 1978.
- Xie, Danyang. Power risk aversion utility functions. *Annals of Economics and Finance*, 1(2): 265–282, 2000.

A Derivation of the First-Order Conditions and $\epsilon^{e,f}$

The Lagrangian underpinning (1) is

$$\begin{aligned}
 L \equiv & u(c_1^f) + \beta u(CE(\tilde{c}_2^f)) \\
 & + \mu_1(y_1^f + y_1^e - c_1^f - s_1^f - s_1^e) \\
 & + \mu_2(\tilde{y}_2^f + \tilde{y}_2^e + s_1^f R_2^f + s_1^e R_2^e - \tilde{c}_2^f) \\
 & + \lambda_1 c_1^f \\
 & + \lambda_2 \tilde{c}_2^f \\
 & + \lambda^f (s_1^f - \tilde{y}_2^f) \\
 & + \nu^f (y_1^f - s_1^f) \\
 & + \lambda^e s_1^e \\
 & + \nu^e (y_1^e - s_1^e)
 \end{aligned}$$

The choice variables are c_1^f , \tilde{c}_2^f , s_1^f , and s_1^e . Instead of relying on Kuhn-Tucker shortcuts, L spells out every Lagrange multiplier. We will shortly collapse the choice variables to just s_1^f and s_1^e , but starting from this expanded perspective illustrates why that simplification is warranted.

L 's first-order conditions are

$$\begin{aligned}
c_1^f &: u'(c_1^f) - \mu_1 + \lambda_1 &= 0 \\
c_2^f &: \beta u'(CE(\tilde{c}_2^f)) CE'(\tilde{c}_2^f) - \mu_2 + \lambda_2 &= 0 \\
s_1^f &: -\mu_1 + \mu_2 R_2^f + \lambda^f - v^f &= 0 \\
s_1^e &: -\mu_1 + \mu_2 R_2^e + \lambda^e - v^e &= 0 \\
\mu_1 &: y_1^f + y_1^e - c_1^f - s_1^f - s_1^e &\geq 0 \\
\mu_2 &: \tilde{y}_2^f + \tilde{y}_2^e + s_1^e R_2^e + s_1^f R_2^f - \tilde{c}_2^f &\geq 0 \\
\lambda_1 &: c_1^f &\geq 0 \\
\lambda_2 &: \tilde{c}_2^f &\geq 0 \\
\lambda^f &: s_1^f - \tilde{y}_2^f &\geq 0 \\
v^f &: y_1^f - s_1^f &\geq 0 \\
\lambda^e &: s_1^e &\geq 0 \\
v^e &: y_1^e - s_1^e &\geq 0
\end{aligned}$$

The properties of u and ψ that satisfy L 's second-order conditions (Gollier, 2001) also guarantee $c_1^{f*}, \tilde{c}_2^{f*} > 0$, and that the two resource constraints hold with equality. As a consequence, $\mu_1^{f*}, \mu_2^{f*} > 0$ and $\lambda_1^*, \lambda_2^* = 0$. Because the consumption path is then fully determined by state variables and saving amounts, c_1^f and \tilde{c}_2^f become redundant choice variables.

Applying those simplifications yields the shorter Lagrangian

$$\begin{aligned}
\mathcal{L} \equiv & u \left(y_1^f + y_1^e - s_1^f - s_1^e \right) + \beta u \left(CE \left(\tilde{y}_2^f + \tilde{y}_2^e + s_1^f R_2^f + s_1^e R_2^e \right) \right) \\
& + \lambda^f \left(s_1^f - y_2^f \right) \\
& + v^f \left(y_1^f - s_1^f \right) \\
& + \lambda^e s_1^e \\
& + v^e \left(y_1^e - s_1^e \right)
\end{aligned}$$

with first-order conditions

$$\begin{aligned}
s_1^f : & -u' \left(c_1^f \right) + \beta u' \left(\tilde{c}_2^f \right) CE' \left(\tilde{c}_2^f \right) R_2^f + \lambda^f - v^f & = 0 \\
s_1^e : & -u' \left(c_1^e \right) + \beta u' \left(\tilde{c}_2^e \right) CE' \left(\tilde{c}_2^e \right) R_2^e + \lambda^e - v^e & = 0 \\
\lambda^f : & s_1^f - \tilde{y}_2^f & \geq 0 \\
v^f : & y_1^f - s_1^f & \geq 0 \\
\lambda^e : & s_1^e & \geq 0 \\
v^e : & y_1^e - s_1^e & \geq 0
\end{aligned}$$

Conveniently, \mathcal{L} recasts L 's consumption-saving tradeoffs in terms of saving decisions alone. This is the mathematical motivation for (1).

The Lagrange multipliers on those saving variables cannot be characterized in general. However, because experimental incentives are usually smaller than field incentives, s_1^{f*} should be interior to lifecycle field income in most cases. That would imply $\lambda^{f*}, v^{f*} = 0$. But, one can certainly think of circumstances where that assumption might not hold (e.g., an experiment conducted in a developing country where y_1^e is on the scale of several months' field income).

To illustrate how this setup admits boundary values of s_1^e , Figure 8 contains two related

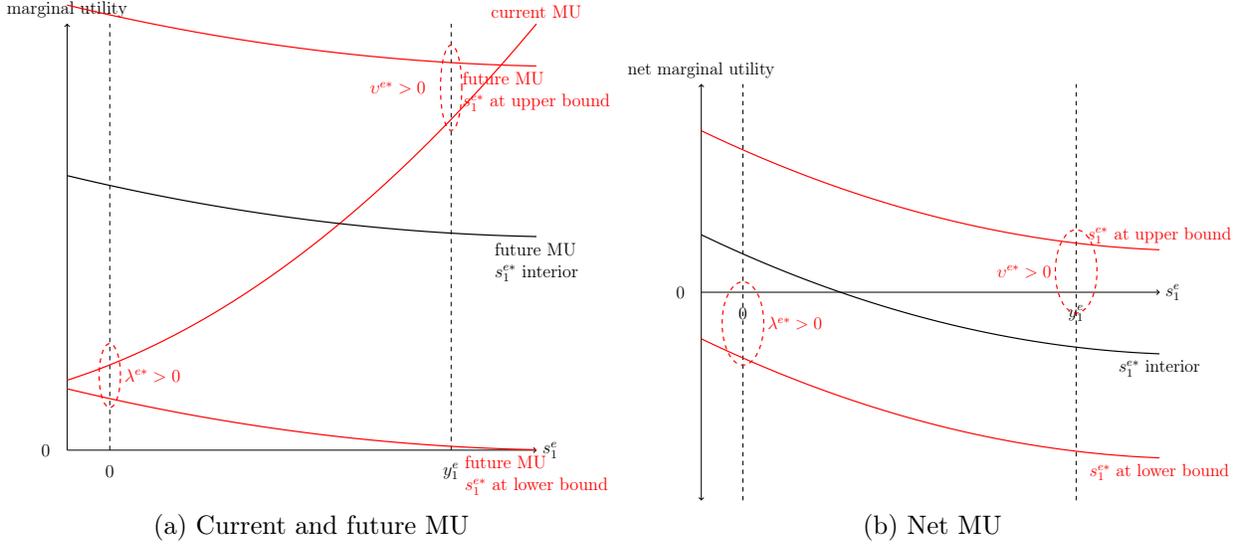


Figure 8: Two visualizations of λ^e and v^e using the marginal utilities from s_1^e 's first-order condition

visualizations of the Lagrange multipliers λ^e and v^e . The first is rooted in an intuitive decomposition of s_1^e 's first-order condition into its current and future marginal utility (MU) components, which is plotted as a supply-and-demand system (Carroll and Kimball, 2005; Bostian and Heinzl, 2018). Under this reading, the participant's future self demands resources via saving, which its current self supplies. The second is the first-order condition itself: given the supply-and-demand interpretation, this equation can be viewed as the net MU ("consumer surplus") from the participant's mental trade.

As these graphs show, interior values of s_1^{e*} occur when the participant's current self can exactly supply the saving that its future self demands (equivalently, when net MU is 0). But, if current MU is always higher than future MU within the allowed saving interval $[0, y_1^e]$, the participant chooses $s_1^{e*} = 0$, and the utility gap $\lambda^{e*} > 0$ appears on the left boundary. Similarly, if future MU is always higher than current MU within that interval, the participant chooses $s_1^{e*} = y_1^e$, and the utility gap $v^{e*} > 0$ appears on the right boundary.

Figure 8 illustrates these gaps by shifting future MU alone. Such shifts could be caused by changes to y_2^e , y_2^f , R_2^e or R_2^f . Of course, these rather simplistic shifts are not the only way to create boundary decisions. As a rule, any experimental incentives that are too stingy

relative to the field will result in $s_1^{e*} = 0$, and any that are too rich will result in $s_1^{e*} = y_1^e$.

The elasticity (4) of s_1^f with respect to s_1^e can be developed from the usual definition

$$\epsilon^{e,f} \equiv \frac{d \ln(s_1^f)}{d \ln(s_1^e)} = \frac{ds_1^f}{ds_1^e} \cdot \frac{s_1^e}{s_1^f}$$

The derivative ds_1^f/ds_1^e can be extracted from \mathcal{L} 's first-order conditions. Summing the s_1^f and s_1^e conditions yields

$$\mathcal{O} \equiv -u'(c_1^f) + \beta u'(\tilde{c}_2^f) CE'(\tilde{c}_2^f) \frac{1}{2} (R_2^f + R_2^e) + \frac{1}{2} (\lambda^f - v^f) + \frac{1}{2} (\lambda^e - v^e) = 0$$

The derivative follows by applying the implicit function theorem to \mathcal{O} :

$$\frac{ds_1^f}{ds_1^e} = -\frac{\partial \mathcal{O} / \partial s_1^e}{\partial \mathcal{O} / \partial s_1^f}$$

where

$$\begin{aligned} \frac{\partial \mathcal{O}}{\partial s_1^f} &= u''(c_1^f) + \beta \left[u''(\tilde{c}_2^f) CE'(\tilde{c}_2^f)^2 + u'(\tilde{c}_2^f) CE''(\tilde{c}_2^f) \right] \cdot \frac{1}{2} (R_2^f + R_2^e) R_2^f \\ \frac{\partial \mathcal{O}}{\partial s_1^e} &= u''(c_1^e) + \beta \left[u''(\tilde{c}_2^e) CE'(\tilde{c}_2^e)^2 + u'(\tilde{c}_2^e) CE''(\tilde{c}_2^e) \right] \cdot \frac{1}{2} (R_2^f + R_2^e) R_2^e \end{aligned}$$