

Self Control, Risk Aversion, and the Allais Paradox

Drew Fudenberg* and David K. Levine**

First Version: May 12, 2006

This Version: February 5, 2010

This paper develops a dual-self model that is compatible with modern dynamic macroeconomic theory and evidence, and shows how it leads to a wide range of behavioral anomalies concerning risk, including the Allais paradox. We calibrate the model to obtain a quantitative fit, by extending the simpler “nightclub” model of Fudenberg and Levine [2006] by introducing consumption commitment. We find that most of the data can be explained with subjective interest rates in the range of 1-7%, short-run relative risk aversion of about 2, and a time horizon of one day for the short-run self.

* Department of Economics, Harvard University

** Department of Economics, Washington University in St. Louis

We thank Daniel Benjamin and Jesse Shapiro for helpful comments and a very careful reading of an early draft, and Eduardo Azevedo, Tao Jin, and Lauren Merrill for exceptional research assistance. We are also grateful to Juan Carillo, Ed Glaeser, Glenn Harrison, Yoram Halevy, John Kagel, Stefan Krasa, Franco Peracchi, Drazen Prelec, and seminar participants at Harvard, Ohio State, the University of Illinois Champagne Urbana and the 2006 CEPR conference on Behavioral Economics. NSF grants SES-0426199, SES-0646816, and SES-0851315 provided financial assistance.

1. Introduction

The dual self model sees decisions as involving an intrinsic conflict between the desire for immediate and safe gratification on the one hand, and riskier long term benefits on the other. This conflict is resolved through a decision criterion that puts weights on both desires. A key element of the theory is that this weighting is endogenous and many different alterations in circumstances should have a similar impact on decision making. It predicts “paradoxes” that include those of small stakes risk aversion, the common-consequence and common-ratio versions of the Allais paradox, and preference reversals for delayed rewards. It also predicts lesser known behavioral riddles, including the effect of reducing the probability of reward on preference reversals induced by delayed rewards, the consequence of delaying rewards for Allais-type preference reversals, preferences for randomization and reversals due to cognitive load.

Our earlier [2006] paper made some of these points, and argued that there is evidence that the cost of self-control is convex and not linear. Our first purpose here is to explore the qualitative implications of convex control costs. Our second goal is to obtain a quantitative fit using a model that is compatible with modern dynamic macroeconomic theory and evidence. To do this, we extend the simpler “nightclub” dual-self model of Fudenberg and Levine [2006] by introducing an additional choice of consumption commitment.

A key motivation for the paper is the idea that models in behavioral economics, as in other areas of economics, should do more than simply organize the data from a given experiment, and that it is much better to have a small number of models that explain a large number of facts than the reverse. By the same token, a model is more useful if its parameters can be held constant across a wider range of settings. There are several competing explanations for each of the facts that our model fits; our quantitative analysis uses a single model with a stable set of parameters to explain many facts.

The idea of the model is that agents use cash on hand as a commitment device, so that on the margin they will consume all of any small unexpected winnings. However, when agents win large amounts, they choose to exercise self-control and save some of their winnings. The resulting intertemporal smoothing make the agents less risk averse, so that they are less risk averse to large gambles than to small ones. As we argued previously, existing data on cognitive load suggests that the cost of self-control is not

linear but rather strictly convex.¹ The combination of these two effects: less risk aversion in the long run, and increasing marginal cost of self control underlie most of our results. In particular, they underlie our finding that a convex cost of self-control can generate an Allais paradox. To begin, we derive some qualitative predictions by applying our [2006] model to a range of new situations. One of these is the interaction of uncertainty and delay: We noted in our previous paper that the dual-self model explains the classic time-inconsistent preferences that have been used to motivate the assumption of quasi-hyperbolic discounting. When there are increasing marginal costs of self-control, there is an additional implication: Preference reversal is less likely when the probability of rewards is smaller. This prediction is borne out in the data of Keren and Roelofsma [1995], which we only learned about after writing the first paper.

We then explain how a convex cost of self control can explain both the classic Allais paradox (where the independence axiom is violated with respect to mixing in a common consequence) and the “common ratio” version of the paradox. In the Allais paradox there are two scenarios, each involving two options. Under expected utility theory, the same option must be chosen in each scenario, but in practice people choose different options in the two scenarios. A key element of the paradox is that one of the scenarios involves a much smaller probability of winning a prize. That means that there is less temptation to the short-run self. The reason that our model predicts the Allais paradox is that the convexity of the cost function leads to a particular sort of violation of the independence axiom: Agents should be “more rational” about choices that are likely to be payoff-irrelevant. This is exactly the nature of the violation of the independence axiom in the Allais paradox. Importantly our theory does not explain all possible violations of the independence axiom: If the choices in each of the two Allais scenarios

¹ The work of Baumeister and collaborators (for example, Muraven et al [1998,2000], Galiot et al [2008]) argues that self-control is a limited resource, moreover one that may be measured by blood glucose levels. The stylized fact that people often reward themselves in one domain (for example, food) when exerting more self-control in another (for example, work) has the same implication. This is backed up by evidence from Shiv and Fedorikhin [1999] and Ward and Mann [2000] showing that agents under cognitive load exercise less self-control, for example, by eating more deserts. The first two observations fit naturally with the idea that a common “self-control function” controls many nearly simultaneous choices. The third fits naturally with the hypothesis that self-control and some other forms of mental activity draw on related mental systems or resources. Benahib and Bisin [2005], Bernheim and Rangel [2004], Brocas and Carillo [2005], Loewenstein and O’Donoghue [2005] and Ozdenoren et al [2006] present similar dual-self models, but they do not derive them from a game the way we do, and they do not discuss risk aversion, cognitive load, or the possibility of convex costs of self-control.

were reversed, the independence axiom would still be violated, but our explanation would not apply.

An additional prediction of our model is that the paradoxical choices in these problems should disappear when the payoff is outside the time horizon of the short-run self. There is very little data on what happens to choices in common-consequence Allais problems when there is delay, but our model is consistent with the data of Baucells and Heukamp [2009] on the effect of delay on the common-ratio problem.

Despite the qualitative success of the base model, attempts to calibrate it to laboratory data using logarithmic per-period utility yields implausible parameter estimates.² This leads us to adopt a nested CES/logarithmic specification that preserves logarithmic preferences for long-run portfolio balancing, while allowing a greater level of risk aversion in the laboratory. Specifically, each morning the agent chooses the venue at which the day’s short-term expenditures will be made based on the anticipated spending level. The envelope of this curve is logarithmic, but once the venue is fixed the agent has CES preferences in the short run. This reflects the idea that over a short period of time, the set of things on which the short-run self can spend money is limited, so the marginal utility of consumption decreases fairly rapidly and risk aversion is quite pronounced. Over a longer time frame there are more possible ways to adjust consumption, and also to learn how to use or enjoy goods that have not been consumed before, so that the long-run utility possibilities are the upper envelope of the family of short-run utilities. With the preferences that we specify in this paper, this upper envelope, and thus the agent’s preferences over steady state consumption levels, reduces to the logarithmic form we used in our previous paper.

After developing the theory of “endogenous nightclubs,” we then calibrate it to data on three different paradoxes. Specifically, we analyze Rabin paradox data from Holt and Laury [2002], the Kahneman and Tversky [1979] and Allais versions of the Allais

² Specifically, while our earlier model can explain the examples in Rabin [2000], those examples (such as rejecting a bet that had equal probability of winning \$105 or losing \$100) understate the degree of risk aversion in small-stakes experiments, where agents are risk averse over much smaller gambles, and fitting our earlier model to these small gambles requires parameter values that conflict both with intuition and with other data. Since the first version of this paper was written, Cox et al [2007] conducted a series of experiments to test various utility theories using relatively high stakes. They also observe that the simple logarithmic model is inconsistent with observed risk aversion, and they argue that the simple linear-logarithmic self-control model does not plausibly explain their data. We will be interested to see whether their data is consistent with the more complex model developed here.

paradox, and the experimental results of Benjamin, Brown, and Shapiro [2006], who find that exposing subjects to cognitive load increases their small-stakes risk aversion.³

Our procedure is to find a set of sensible values of the key parameters, namely the subjective interest rate, income, the degree of short-term risk aversion, the time horizon of the short-run self, and the degree of self-control, using a variety of external sources of data. We then investigate how well we can explain the paradoxes using the calibrated parameter values and the dual-self model. How broad a set of parameter values in the calibrated range will explain the paradoxes? To what extent can the same set of parameter values simultaneously explain all the paradoxes? Roughly speaking, we can explain most of the data if we assume an annual subjective interest rate in the range of 1-7%, a short-run “relative risk aversion of consumption” of about 2, and a daily time horizon for the short-run self. We find that the Rabin paradox is relatively insensitive to the exact parameters assumed; the Allais paradox is sensitive to choosing a plausible level of risk aversion; and the Chilean cognitive load data is very sensitive to the exact parameter values chosen. In particular, the Chilean data requires a subjective interest rate of 7%, which is at the high end of the range consistent with macroeconomic data.

We should emphasize that the subject populations we study are heterogeneous and range from Chilean high school students to U.S. and Dutch undergraduates. Moreover, only some subjects exhibit reversals. In short, there is no reason to believe that the reversals we observe in the data can all be explained by a single common set of “representative” parameter values. Consequently instead of trying to focus on the single parameter constellation that would best fit all the data, we focus on the range and robustness of the parameter values that can be used to explain reversals, and the fact that even across subjects and populations this range is roughly the same.

³ The main focus of Benjamin, Brown and Shapiro [2006], like that of Frederick [2005], is on the correlation between measures of cognitive ability and the phenomena of small-stakes risk aversion and of a preference for immediate rewards. Benjamin, Brown and Shapiro find a significant and substantial correlation between each of these sorts of preferences and cognitive ability. They also note that the correlation between cognitive ability and time preference vanishes when neither choice results in an immediate payoffs, and that the correlation between small-stakes risk aversion and “present bias” drops to zero once they control for cognitive ability. This evidence is consistent with our explanation of the Rabin paradox, as it suggests that that small-stakes risk aversion results from the same self-control problem that leads to a present bias in the timing of rewards. They also discuss the sizable literature that examines the correlation between cognitive ability and present bias without discussing risk aversion.

After showing that the base model provides a plausible description of data on attitudes towards risk, gambles, and cognitive load, we examine the robustness of the theory. Specifically, in the calibrations we assume that the opportunities presented in the experiments are unanticipated, so we consider what happens when gambling opportunities are foreseen.

2. The Base Model

To begin we review the base model from our [2006] paper. We consider an infinite-lived consumer making a savings decision. Each period $t = 1, 2, \dots$ is divided into two sub-periods, the “bank” sub-period and the “nightclub” sub-period; the consumer discounts utility between periods using discount factor δ , but no discounting occurs between the sub-periods. Wealth at the beginning of the bank sub-period is denoted by w_t . During the “bank” sub-period, consumption is not possible, and wealth is divided between savings s_t , which remains in the bank, pocket cash x_t , which is carried to the nightclub. In the nightclub consumption $0 \leq c_t \leq x_t$ is determined, with $x_t - c_t$ returned to the bank at the end of the period. Wealth next period is just $w_{t+1} = R(s_t + x_t - c_t)$, where $R \in (0, \delta)$ is the risk-free interest rate. For simplicity money returned to the bank bears the same rate of interest as money left in the bank.

Our underlying rationale is that perfect capital markets are available at the bank, so we can capitalize all of the consumer’s future income into her initial wealth w_1 . The only constraint at the bank then is that wealth w_t must be non-negative. By way of contrast, capital markets are not available at the nightclub, and no income is received there, so the only choice at the nightclub is how much pocket cash to spend.

The utility of the short-run self living in period t is $u(c_t)$. We assume this is twice continuously differentiable, strictly differentially concave, and satisfies the Inada condition that $\lim_{c_t \rightarrow \infty} u'(c_t) = \infty$. The long-run self maximizes the expected discounted present value of the utility of the short-run selves, subject to a cost of self-control. This self-control cost depends on the resources the short-run self perceives as available to himself, which in turn determine a *temptation utility* for the short-run self, representing the utility the short-run self perceives as available if no self-control is used. Denote this temptation utility by \bar{u}_t . The actual realized utility that the long-run self allows the short-run self is u_t , and there may be cognitive load due to other activities, d_t .

Then the cost of self-control is $g(d_t + \bar{u}_t - u_t)$. The key idea here is that the cost of self-control depends on the difference between the utility by which the short-run self is tempted, \bar{u}_t , and the utility the short-run self is allowed, u_t .

In the bank no consumption is possible, and so there is no temptation for the short-run self. In the nightclub the short-run self cannot borrow, and wishes to spend all of the available pocket cash x_t on consumption. Hence $\bar{u}_t = u(x_t)$.

The problem faced by the long-run self is to choose pocket cash and consumption to maximize the present value using the discount factor δ of short-run self utility net of the cost of self-control. The objective function of the long-run self is

$$U_{RF} = E \sum_{t=1}^{\infty} \delta^{t-1} [u(c_t) - g(u(x_t) - u(c_t))] \quad (2.1)$$

which is to be maximized with respect to $c_t \geq 0, x_t \geq 0$ subject to w_1 given, $w_{t+1} = R(s_t + x_t - c_t)$, $s_t + x_t \leq w_t$ and $w_t \geq 0$. Notice that this is a simple optimization problem with no uncertainty and perfect foresight.

In this formulation there is a single long-run self with time-consistent preferences. Although the impulsive short-run selves are the source of self-control costs, the equilibrium of the game between the long-run self and the sequence of short-run selves is equivalent to the optimization of this reduced-form control problem by the single long-run self. A crucial aspect of the model is the ‘‘pocket cash’’ x_t that serves to ration consumption and so reduce the temptation to the short-run self. In the simple perfect-foresight version of the model, it will be optimal to give the short-run self exactly the amount to be spent at the nightclub, and so avoid temptation and self-control cost entirely. In effect the long-run self hands the pocket cash to the short-run self to take to the nightclub and says ‘‘here...go crazy....spend it all.’’ The actual decision about how much pocket cash to allocate to the short-run self is taken at a location – ‘‘the bank’’ – where there are no tempting consumption possibilities.

In our earlier paper the notion of a bank and pocket cash were taken literally. In practice there are many strategies that individuals use to reduce the temptation for impulsive expenditures. The view we take here is that pocket cash is determined by mental accounting of the type discussed by Thaler [1980], and not necessarily by physically isolating money in a bank. In other words, x_t should not be viewed as the literal amount of money the short-run self has in their wallet or the amount available

including cash cards and so forth, but should be viewed as the amount of resources that the short-run self feels entitled to use. The strategies individuals use for this type of commitment can be varied. For example some people may choose to carry only a limited amount of cash and no credit cards. Others may allow the short-run self to spend money only in the “right pocket.” Yet others may engage in more direct mental accounting of the form “you may spend \$100 at the nightclub, but no more.” Although this sort of mental account is not directly observable, we can calculate x_t from knowledge of the underlying parameters of preferences, and to the extent that those are stable across situations the “cash budget” will be too.

The base model is easy to solve. Because there is no cost of self-control in the bank, the solution to this problem is to choose $c_t = x_t$. In other words, cash x_t is chosen to equal the optimal consumption for an agent without self-control costs. The agent then spends all pocket cash at the nightclub, and so incurs no self-control cost there. Since $c_t = x_t$, the utility of the short-run self is $u(x_t)$, and as there is no self-control cost, this boils down to maximizing

$$\sum_{t=1}^{\infty} \delta^{t-1} u(x_t)$$

subject to the budget constraint $w_{t+1} = R(w_t - x_t)$. Denote the solution to this problem as \hat{x}_t , and the corresponding value function $V(w_t)$.

3. Risky Drinking: Nightclubs and Lotteries

Suppose in period 1 (only) that when the agent arrives at the nightclub of her choice, she has the choice between two lotteries, S and L , with intertemporal returns $\tilde{z}_t^S, \tilde{z}_t^L |_{t=1}^{\infty}$. We will adopt the convention that when there is a high self-control cost preferences for the lottery S are stronger (relative to L) than when there is no self-control cost. Intuitively we think of S as better liked by the short-run self and L by the long-run self. In our applications the random receipts will be independent between periods. Initially we will assume that this choice is completely unanticipated – that is, it has prior probability zero. What is the optimal choice of lottery given x_1 ? For simplicity, we assume throughout that the agent does not expect to have more choices between lotteries at nightclubs after period 1. We should emphasize that this is for convenience only: the overall savings and utility decision will not change significantly provided that

the probability of getting future choice opportunities is small. Adding this small probability explicitly will add a great many complicated but essentially irrelevant terms to our equations. We examine the implications of anticipated choice in more detail below.

The returns $\tilde{z}_t^S, \tilde{z}_t^L \big|_{t=1}^\infty$ may be positive or negative, but we suppose that the largest possible loss in period 1 is less than the agent's pocket cash. There are number of different ways that the dual-self model can be applied to this setting, depending on the timing and "temptingness" of the choice of lottery and spending of its proceeds. In this paper we assume that the short-run self in the nightclub simultaneously decides which lottery to pick and how to spend in the first period for each possible realization of the lottery.

Since the highest possible short-run utility comes from consuming the entire proceeds of the lottery, the temptation utility is $\max\{Eu(x_1 + \tilde{z}_1^S), Eu(x_1 + \tilde{z}_1^L)\}$ where \tilde{z}_1^j is the realization of lottery $j = S, L$. This temptation must be compared to the expected short-run utility from the chosen lottery. If we let $\tilde{c}_1^j(z_1^j)$ be the consumption chosen contingent on the realization of lottery j , the self-control cost is

$$g(\max\{Eu(x_1 + \tilde{z}_1^S), Eu(x_1 + \tilde{z}_1^L)\} - Eu(\tilde{c}_1^j)).$$

Let V denote the maximized value of LR's discounted expected payoff, as a function of wealth. (From the stationarity of the problem this value does not depend on t .) Then the decision problem of the long-run self in period 1 is to choose $\tilde{c}_1^j(z_1^j)$ and either S or L to maximize

$$\begin{aligned} & Eu(\tilde{c}_1^j) - g(\max\{Eu(x_1 + \tilde{z}_1^S), Eu(x_1 + \tilde{z}_1^L)\} - Eu(\tilde{c}_1^j)) \\ & + \delta EV(R(w_1 + \tilde{z}_1^j - \tilde{c}_1^j) + \sum_{t=2}^\infty R^{2-t} \tilde{z}_t^j) \end{aligned}$$

In what follows, it will be useful to define

$$w_2^j(c_1^j) \equiv R(w_1 + \tilde{z}_1^j - c_1^j) + \sum_{t=2}^\infty R^{2-t} \tilde{z}_t^j.$$

Consider first the problem of determining the optimal consumption in period 1 as a function of the income receipt z_1^j . Let γ_1 denote the marginal cost of self-control in the first period. We let $\hat{c}_1^j(\gamma_1)(z_1^j)$ be the solution to the first-order condition for a maximum for a given marginal cost of self-control, that is, the unique solution to

$$(1 + \gamma_1)u'(c_1^j) = \delta R EV'(w_2^j(c_1^j)),$$

and find the corresponding marginal cost of self-control

$$\hat{\gamma}_1^j(\gamma_1) = g'(\max\{Eu(x_1 + \tilde{z}_1^S), Eu(x_1 + \tilde{z}_1^L)\} - Eu(\min\{\hat{c}_1^j(\gamma_1)(\tilde{z}_1^j), x_1 + \tilde{z}_1^j\}))$$

Theorem 1: For given x_1 and each $j \in \{S, L\}$ there is a unique solution to

$$\gamma_1^j = \hat{\gamma}_1^j(\gamma_1^j)$$

and this solution together with $\tilde{c}_1^j = \min(\hat{c}_1^j(\gamma_1^j)(z_1^j), x_1 + z_1^j)$ and the choice of j that maximizes long-run utility is necessary and sufficient for an optimal solution to the consumer's choice between lotteries S and L .




Figure 1 - The "Consumption" Function

The “consumption function” is $\tilde{c}_1^j = \min(\hat{c}_1^j(\gamma_1^j)(z_1^j), x_1 + z_1^j)$. Let \hat{z} such that all the winnings are spent, that is, $\hat{c}_1^j(\gamma_1)(\hat{z}) = x_1 + \hat{z}$. Note that for arbitrary values of x_1 we may have \hat{z}_1 negative. The function \hat{c}_1^j is sketched in Figure 1. For $z_1^j < \hat{z}_1$ no self-control is used, and all winnings are spent. Above this level self-control is used, with only a fraction of winnings consumed, and the rest going to savings. When the time period is short, \hat{c}_1^j is very flat, so that only a tiny fraction of the winnings are consumed immediately when receipts exceed the critical level. Thus when the agent is patient he is almost risk neutral with respect to large gambles. However the agent is still risk averse to small gambles, as these will not be smoothed but will lead to a one for one change in current consumption.⁴

4. Qualitative Analysis

We now turn to a qualitative analysis of the choice between lotteries. We consider a variety of conceptual experiments involving both one-period and intertemporal choices.

Uncertainty and Time Delay: We noted in our previous paper that the dual-self model explains the classic time-inconsistent preferences that have been used to motivate the assumption of quasi-hyperbolic discounting. When there are increasing marginal costs of self-control, there is an additional implication: Preference reversal is less likely when the probability of rewards is smaller.

Specifically, suppose that unexpectedly in period 1 the short-run self at the nightclub is offered a choice between an amount z_t^S in period t and an amount $z_{t+1}^L = \theta z_t^S$ in period $t + 1$ where $\theta > R$. If period t is in the future, that is, $t > 1$ then since no self control costs are involved and $\theta > R$ option L is strictly preferred.

The situation is different if $t = 1$. Notice that for small amounts $z_1^S < \hat{z}_1$, if the amount today is taken it is consumed immediately. Hence the utility gain from choosing S is

⁴ Note that the non-differentiability of consumption and thus of indirect utility in our model occurs at the self-control threshold, which is strictly positive and is endogenous to the model (but exogenous to the specific gamble). Thus the non-differentiability is different than the kink at zero in loss-aversion models, and also different than kinks at “reference points” that vary with the gamble under consideration.

$$\begin{aligned} \Delta &= u(x_1 + z_1^S) + \delta EV(w_2) \\ &\quad - u(x_1) - g(u(x_1 + z_1^S) - u(x_1)) + \delta V(w_2 + \theta z_1^S) \end{aligned}$$

Differentiating this with respect to z_1^S we find

$$d\Delta / dz_1^S = (1 - g'(0))u'(x_1) + \delta\theta V'(w_2).$$

As is standard the first-order condition for the value function without self-control that defines V implies that $V'(w_2) = (1/R\delta)u'(x_1)$. Hence if $g'(0) > 0$ then S is the preferred alternative for small z_1^S .

This shows that preferences can exhibit a reversal in the sense that the willingness to wait in period 1 for higher payoff in period 2 is different than willingness to postpone payoffs from period $t > 1$ to $t + 1$. The fact that some subjects exhibit such a reversal has been used to motivate the assumption of quasi-hyperbolic utility. In our model, this preference reversal arises because at the nightclub the short-run self is rationed by the available cash x_1 . An extra cash payment of z_1 today will cause a temptation to increase spending that is costly for the long-run self to control. By contrast there is no temptation associated with future payoffs, and so there can be “preference reversal” whenever the cost of resisting the short-run temptation is sufficiently high. Data from Keren and [1995] in Table 1 in which the “Probability of reward” is equal to 1.0 shows how this happens in practice.⁵

Table 1 – Dynamic Preference Reversal

		Probability of reward ⁶	
Scenario		1.0 (60)	0.5 (100)
1	S \$175 now	0.82	0.39
	L \$192 4 weeks	0.18	0.61
2	S \$175 26 weeks	0.37	0.33
	L \$192 30 weeks	0.63	0.67

⁵ This experimental result is confirmed by Weber and Chapman [2005], and discussed in Halevy [2008], who proposed an objective function that is consistent with these choices. Note that the experiment was in Dutch Florins. We converted from Dutch Florins to U.S. Dollars using an exchange rate typical of the early 1990s of 1.75 Florin per Dollar.

⁶ Sample size in parentheses.

Here 82% of the population will take the smaller but more immediate reward, when the earlier reward is “now” but only 37% will do so when the earlier reward is not for 26 weeks.

However, the self-control model with increasing marginal cost of self-control has a second implication. Suppose that instead of a certain reward z_t^S, z_{t+1}^L there is only a chance p of getting the reward. When $t > 1$ standard expected utility theory applies, and this makes no difference: L is still strictly preferred as long as $\theta > R$. Consider however the utility difference for $t = 1$. Still assuming $z_1^S < \hat{z}_1$ so that $c_1^S = x_1 + z_1^S$ we have

$$\begin{aligned} \Delta(p) &= p[u(c_1^S) + \delta V(R(w_1 + z_1^S - c_1^S))] \\ &\quad - p[u(x_1) - (1/p)g(p[u(x_1 + z_1^S) - u(x_1)])] + \delta V(R(w_1 - c_1^L) + \theta z_1^S) \end{aligned}$$

From this we may compute

$$\begin{aligned} \Delta'(1) &= \Delta(1) \\ &\quad + g'(u(x_1 + z_1^S) - u(x_1))[u(x_1 + z_1^S) - u(x_1)] - g([u(x_1 + z_1^S) - u(x_1)]) \\ &> \Delta(1) \end{aligned}$$

where the final inequality holds strictly if $g'(\cdot)$ is strictly increasing. In particular, if $\Delta(1) = 0$ so there is exact indifference when $p = 1$ then a small decrease in p will imply $\Delta(p) < 0$, so that L will be strictly preferred. By continuity, if $\Delta(1)$ is slightly positive so that A is preferred, then there will be $p < 1$ such that L is preferred. Or put differently, as p is reduced, behavior at $t = 1$ comes to resemble that for $t > 1$. This dependence of the choices on the probability of reward is not consistent with quasi-hyperbolic preferences (as in Laibson [1997]) or with the version of the independence axiom (for choices over menus) imposed as an axiom by Gul and Pesendorfer [2001] and Dekel, Lipman, and Rustichini [2008]. It is however in accord with the Keren and Roelofsma [1995] data in Table 1, where when p is reduced from one to 0.5 in both scenario 1 and scenario 2 the probability of choosing the early reward S is quite similar, 0.39 and 0.33 respectively, and also very similar to the probability of choosing the early reward when $p = 1$ and $t > 1$, which is 0.37.

Risk Aversion in the Short-Run and the Long-Run: We start by establishing a basic relationship between long-run risk aversion as measured by the value function and short-term risk aversion as measured by the utility function

Theorem 2: Let $\xi \equiv \sup_{c_1} (-c_1 u''(c_1)/u'(c_1))/\inf_{c_1} (-c_1 u''(c_1)/u'(c_1))$. Then

$$-\frac{V''(w_1)}{V'(w_1)} \leq -\xi \frac{c_1}{w_1} \frac{u''(c_1)}{u'(c_1)}$$

Proof: See the Appendix.

In this model period length is the time-horizon of the short-run self, which is on the order of a day. The ratio of daily consumption to lifetime wealth is quite small, on the order of 1/10000. This means that consumption that is spread over the lifetime represents four orders less magnitude of risk than consumption that is taken in the initial period only. In this model, with positive self-control, we have already observed that unanticipated gains less than \hat{z}_1 are all consumed immediately. This means they are evaluated using u and represent a great risk. On the other hand, larger or anticipated gains are spread over the entire lifetime and represent a much less great risk. As we showed in Fudenberg and Levine [2006] this predicts the small-stakes risk aversion known as the Rabin paradox. For example, if we observe risk aversion in the laboratory and compute risk aversion with respect to lifetime wealth, we get numbers for the coefficient of relative risk aversion on the order of 8000, while stock market behavior reveals relative risk aversion of around 8. Macroeconomic studies often take the coefficient of relative risk aversion to be one (that is, logarithmic utility). The resulting four orders of magnitude difference are roughly what we would expect if the laboratory winnings are smaller than \hat{z}_1 , and stock market gains and losses larger. We will explore this quantitatively in later sections. This differential between risk aversion for small and large gains also plays a key role in our analysis of reversals such as the Allais and common ratio paradoxes.

The Allais and Common-Ratio Paradoxes: We turn now to the analysis of period one gambles, so that $\tilde{z}_t^j = 0, t > 1$. For simplicity we consider three possible outcomes $Z_b < Z_m < Z_g$ where Z_b is taken to be zero.

We consider two different scenarios. In the base scenario the probabilities of the outcomes $i \in \{b, m, g\}$ are p_i^j under the alternatives $j = S, L$. In the second scenario these probabilities are \bar{p}_i^j .

In the Allais paradox the probability of the good outcome is the same in both scenarios $\bar{p}_g^j = p_g^j$, while the probability of the middle outcome is reduced by the same amount for both actions $\bar{p}_m^j = p_m^j - P$. For example, in the classical Allais paradox, the rewards are $Z_b = 0, Z_m = 1000000, Z_g = 5000000$, while the base probabilities are $p^S = (0, 1, 0)$, meaning the middle outcome is certain, and $p^L = (.01, .89, .10)$. The second scenario is defined by $P = .89$, meaning that $\bar{p}^S = (0, .11, 0)$ and $\bar{p}^L = (0, 0, .10)$. An expected utility maximizer must make the same choice in each scenario, but it often observed that when asked about these hypothetical choices, many subjects prefer S in the base scenario and L in the second scenario.

In the common-ratio paradox, the probability of both the middle and good outcome are reduced by a common factor α , so that $\bar{p}_M^j = \alpha p_M^j, \bar{p}_G^j = \alpha \bar{p}_G^j$. In Baucells, Heukamp and Villasis [2007], for example, $Z_b = 0, Z_m = 9, Z_g = 12$. The base probabilities are $p^S = (0, 1, 0), p^L = (0, 0, .8)$, and $\alpha = 0.1$ so that $\bar{p}^S = (.9, .1, 0), \bar{p}^L = (.92, 0, .08)$. As with the Allais paradox, an expected utility maximizer makes the same choice in each scenario, yet in practice many choose S in the base scenario and switch to L in the second scenario. For example Baucells, Heukamp and Villasis [2007] data shown in Table 2 show that only 22% choose the safer choice S in the second scenario, while the substantially greater fraction 58% make that choice in the base scenario.⁷

Table 2: Common Ratio Paradox

S 1.00 chance of 9 euros	0.58
L 0.80 chance of 12 euros	
S 0.10 chance of 9 euros	0.22
L 0.08 chance of 12 euros	

⁷ Based on a sample of 221 subjects.

To apply the dual-self model, we compute the utility advantage of S .

$$\begin{aligned} \Delta = & \{Eu(\tilde{c}_1^S) - g(Eu(x_1 + \tilde{z}_1^S) - Eu(\tilde{c}_1^S)) + \delta EV(w_2(\tilde{c}_1^S))\} \\ & - \{Eu(\tilde{c}_1^L) - g(Eu(x_1 + \tilde{z}_1^S) - Eu(\tilde{c}_1^L)) + \delta EV(w_2(\tilde{c}_1^L))\} \end{aligned}$$

Notice that we have assumed that S is a less risky alternative than L . According to Theorem 2 if c_1/w_1 is small then if L is preferred by u to S then it is also preferred by the less risk averse V , and regardless of self-control, L is the optimal choice and no reversal will occur. The interesting case, therefore, is the case in which $Eu(x_1 + \tilde{z}_1^S) \geq Eu(x_1 + \tilde{z}_1^L)$, so we examine this case.

So that we may use calculus, we take $p_i^L = p_i^S + \lambda P_i$ where $\lambda > 0$ is small, and $P_b = -P_m - P_g$. To determine the optimal choice, we compute $d\Delta/d\lambda$:

$$\begin{aligned} d\Delta/d\lambda = & \\ & [1 + g'(Eu(x_1 + \tilde{z}_1^S) - Eu(\tilde{c}_1^S))] [(u(c_{1m}^S) - u(c_{1b}^S))P_m + (u(c_{1g}^S) - u(c_{1b}^S))P_g] \\ & + \delta [(V(w_2(c_{1m}^S)) - V(w_2(c_{1b}^S)))P_m + (V(w_2(c_{1g}^S)) - V(w_2(c_{1b}^S)))P_g] \end{aligned}$$

If this is positive then for small λ the optimal choice is L , and if it is negative the optimal choice is S . When can changing from the base scenario to the second scenario reverse the sign of $d\Delta/d\lambda$? In the Allais paradox $\bar{p}_g^j = p_g^j$, and $\bar{p}_m^j = p_m^j - P$. In other words $\bar{P}_i = P_i$. In the common ratio paradox $\bar{p}_M^j = \alpha p_M^j$, $\bar{p}_G^j = \alpha p_G^j$ so $\bar{P}_i = \alpha P_i$. Thus when the two choices are close together, the Allais paradox can be analyzed as a special case of the common ratio paradox in which $\alpha = 1$.

If we hold fixed $\gamma_1 = g'(Eu(x_1 + \tilde{z}_1^S) - Eu(\tilde{c}_1^S))$ we see that

$$d\bar{\Delta}/d\lambda = \alpha d\Delta/d\lambda$$

so that there can be no sign change and no reversal. If

$$U^* = (u(c_{1m}^S) - u(c_{1b}^S))P_m + (u(c_{1g}^S) - u(c_{1b}^S))P_g$$

and

$$V^* = (V(w_2(c_{1m}^S)) - V(w_2(c_{1b}^S)))P_m + (V(w_2(c_{1g}^S)) - V(w_2(c_{1b}^S)))P_g$$

have the same sign, then the magnitude of γ_1 will not matter, and again there will be no reversal. Since u is more risk averse than V the interesting case then is one in which U^* is negative (S preferred) and V^* is positive (L preferred).

Let us suppose that $Z_b < \hat{z}_1$. Then

$$\gamma_1 = g'((u(x_1 + z_{1m}^S) - Eu(\tilde{c}_{1m}^S))p_m^S + (u(x_1 + z_{1g}^S) - Eu(\tilde{c}_{1g}^S))p_g^S).$$

In both the Allais paradox and the common ratio paradox $\bar{p}_m^S < p_m^S, \bar{p}_m^L \leq p_m^L$, so if $Z_m > \hat{z}_1$ and $g'(\cdot)$ is strictly increasing $\bar{\gamma}_1 < \gamma_1$. If $d\Delta/d\lambda = 0$, this implies $d\bar{\Delta}/d\lambda < 0$, so by continuity the same is true if $d\Delta/d\lambda$ is slightly positive. In other words, the dual self model with strictly convex cost of self-control predicts both an Allais and common ratio paradox.

Reversal and Delay: Dual self theory has a second implication for Allais and common ratio paradoxes: if the results of gambles are delayed long-enough that they fall outside the time horizon of the short-run self, the reversals should disappear. Some people find this conclusion to be implausible. There is very little data on what happens to Allais type reversals when there is delay, but Baucells and Heukamp [2009] have studied the common ratio reversal with and without delay. Their data is summarized in Table 3..

Table 3: Common Ratio with Delay

	Now	3 month delay
S 1.00 chance of 9 euros	0.58	0.43
L 0.80 chance of 12 euros		
S 0.10 chance of 9 euros	0.22	0.21
L 0.08 chance of 12 euros		

The key fact is that the reversal rate is 0.36 without delay and falls to 0.22 with delay – exactly as the theory predicts.

Betweenness: Suppose that in addition to S and L , there is a third option M consisting of a probability $0 < \lambda < 1$ of S and $1 - \lambda$ of L . In expected utility theory unless S and L are indifferent, M can never be chosen, and indeed, can never be strictly

preferred over both S and L . Dekel [1986] proposed this “betweenness” property as a way of relaxing the independence axiom. However, Camerer and Ho [1994] summarize considerable past work showing that this axiom is violated by a non-trivial fraction of subjects, and also conducted their own experiments that yielded similar results. For example, 20% of their subjects express a strict preference for the intermediate choice M .⁸

From the perspective of our theory this is not so surprising. Suppose that S is preferred by u and L by V . If L is chosen there is a high temptation so a high marginal cost of self-control. With this high marginal cost of self-control, S may look quite attractive. If S then there is no temptation so a low marginal cost of self-control. This may make L attractive, suggesting that the optimum may lie in between.

To establish this formally, consider the utility of M as a function of λ

$$\begin{aligned} U(\lambda) &= \lambda Eu(\tilde{c}_1^S) + (1 - \lambda)Eu(\tilde{c}_1^L) \\ &\quad - g\left(Eu(x_1 + \tilde{z}_1^S) - (\lambda Eu(\tilde{c}_1^S) + (1 - \lambda)Eu(\tilde{c}_1^L))\right) \\ &\quad + \delta\left[\lambda EV(w_2(\tilde{c}_1^S)) + (1 - \lambda)EV(w_2(\tilde{c}_1^L))\right] \end{aligned}$$

Differentiating with respect to λ we find

$$U'(\lambda) = (1 + \gamma_1(\lambda))(Eu(\tilde{c}_1^S) - Eu(\tilde{c}_1^L)) + \delta[EV(w_2(\tilde{c}_1^S)) - EV(w_2(\tilde{c}_1^L))]$$

Notice that $\gamma_1(0) < \gamma_1(1)$ provided that g is strictly concave. If $\gamma_1(0)$ is small enough then $U'(0) > 0$, while if $\gamma_1(1)$ is large enough $U'(1) < 0$. These two conditions imply that there is some $0 < \lambda < 1$ such that M is strictly preferred to both S and L .

5. Making the Evening's Plans: Robustness

Our base model supposes that the choice between S and L is completely unanticipated. How does the optimal choice of pocket cash x_1 change if the decision maker realizes that she will face a gamble? Specifically, let π^G denote the probability of

⁸ We are grateful to Simone Cerreia-Vioglio for bringing the Camerer and Ho paper to our attention. They did not solicit willingness to pay for the most preferred option, so while stating indifference was an option, it may be that some subjects were indifferent but chose not to say so. Camerer and Ho conducted different sets of experiments with two populations one of high school subjects and one of MBA students. The most relevant results are those with the MBA students, which is what we report here. Note that sixteen percent of subjects strictly prefer both S and L to the intermediate choice, which can be explained in our theory only with a decreasing marginal cost of self-control.

getting the gambles. Our assumption has been that $\pi^G = 0$. In this case the solution is simple: there is no self-control problem at the bank, so the choice is to spend all the pocket cash in the nightclub.

To examine the robustness of our results, consider then the polar opposite case in which $\pi^G = 1$, that is, the agent knows for certain she will be offered the choice between S and L . Given the choice of pocket cash x_1 the choice of which lottery to choose at the nightclub and how much to spend are the same regardless of the beliefs that led to the choice of x_1 . To keep things simple, we will assume that $\tilde{z}_1^j \ll \hat{z}_1$ so that all the proceeds of the gamble will be spent at the nightclub and x_1 will be chosen to be strictly positive. We specialize to the case of logarithmic preferences $u(c) = \log c$.

In the Appendix we show that

Theorem 3: *First order condition necessary for an optimum are*

$$(1 - \delta)x_1 + \delta \frac{1}{E(x_1 + \tilde{z}_1^j)^{-1}} = (1 - \delta)w_1$$

where j is the chosen alternative S or L .

To understand this condition suppose that \tilde{z}_1^j is constant, not random. Then the first order condition reduces to $x_1 + \delta E\tilde{z}_1^j = (1 - \delta)w_1$. Here $E\tilde{z}_1^j$ does not substitute for pocket cash x_1 on a 1-1 basis, as it has a miniscule effect on life-time wealth, but as δ is nearly one, as we would expect it nearly does so. More generally we can write

$$(1 - \delta)x_1 + \delta E\tilde{z}_1^j + \delta \left[\frac{1}{E(x_1 + \tilde{z}_1^j)^{-1}} - \frac{1}{[E(x_1 + \tilde{z}_1^j)]^{-1}} \right] = (1 - \delta)w_1$$

When the variance of \tilde{z}_1^j is positive $E(x_1 + \tilde{z}_1^j)^{-1} < [E(x_1 + \tilde{z}_1^j)]^{-1}$

so that the term in square brackets is positive. Hence with logarithmic preferences the optimal plan is to hold a little less pocket cash in the face of uncertainty.

6. Quantitative Analysis

We turn now to a quantitative analysis of the reversal phenomenon. Roughly speaking, the question we wish to address is whether there is a single set of preferences

over risk, time, and self-control that can explain the behavior of a median individual across a wide range of decisions.

To make the model operational, we need to extend it in two ways. First, for compatibility with the macro literature and macro data we would like to assume that the period utility is logarithmic. However, data on laboratory experiments indicates that even allowing for the three order of magnitude difference between wealth and daily expenditures subjects are still more risk averse than the logarithm. This leads us to adopt a nested CES/logarithmic specification that preserves logarithmic preferences for long-run portfolio balancing, while allowing a greater level of risk aversion in the laboratory. We do this through the device of consumption commitment – also widely used in the macro literature. Second, to reconcile data with the model, we must recognize that a portion of consumption c_t^d is for durables that is paid for in advance, and so not subject to temptation.⁹

Consumption Commitment: So far, we have followed Fudenberg and Levine [2006]. Now we consider an extension of that model that we will use to explain the degree of risk aversion we observe in experimental data. Specifically, we suppose that there is a choice of nightclubs to go to in the nightclub sub-period. These choices are indexed by the quality of the nightclub $c_t^q \in (0, \infty)$. In a nightclub of quality c_t^q we assume that the utility of the short-run self has the form $u(c_t | c_t^q)$ depending on the amount consumed c_t there and the quality of the nightclub.

The utility function at the nightclub is assumed to satisfy $u(c_t | c_t^q) \leq u(c_t | c_t)$. This means that when planning to consume a given amount c_t it is best to choose the nightclub of the same index. Intuitively, the quality c_t^q of a nightclub represents a “target” level of consumption expenditure at that nightclub. That is, if you are going to consume a low level of c_t you would prefer to spend it at a nightclub with a low value of c_t^q and you would like to consume a high level of c_t at a high quality nightclub.

⁹ Durable and/or committed consumption is a significant fraction (roughly 50%) of total consumption so we need to account for it in calibrating the model, but consumption commitments are not our focus here. For this reason we use a highly stylized model, with consumption commitments reset at the start of each time period. A more realistic model of durable consumption would have commitments that extend for multiple periods, as in Grossman and Laroque [1990].

It is useful to think of a low value of c_t^q as representing a nightclub that serves cheap beer, while a high value of c_t^q represents a nightclub that serves expensive wine. At the beer bar c_t represents expenditure on cheap beer, while at the wine bar it represents the expenditure on expensive wine. The assumption that $u(c_t | c_t^q) \leq u(c_t | c_t)$ captures the idea that spending a large amount at a low quality nightclub results in less utility than spending the same amount at a high quality nightclub: lots of cheap beer is not a good substitute for a nice bottle of wine. Conversely, spending a small amount at a high quality nightclub results in less utility than spending the same amount at a low quality nightclub: a couple of bottles of cheap beer are typically better than a thimble-full of nice wine. People with different income and so different planned consumption levels will choose consumption sites with different characteristics. The quality of a nightclub can also be interpreted as a state variable or capital stock that reflects experience with a given level of consumption: a wine lover who unexpectedly wins a large windfall may take a while both to learn to appreciate differences in *grands crus* and to learn which ones are the best values.¹⁰

We assume that $u(c_t | c_t) = \log c_t$; this ensures that in a deterministic and perfectly foreseen environment without self-control costs, behavior is the same as with standard logarithmic preferences. To avoid uninteresting approximation issues, we assume that there are a continuum of different kinds of nightclubs available, so that there are many intermediate choices between the beer bar and wine bar.

There are a great many possible functional forms satisfying these properties. Our choice of a specification is guided both by analytic convenience and by evidence (examined below) that short-term risk preferences seem more risk averse than consistent with the logarithmic specification, even when self-control costs are taken into account. This leads us to adopt the functional form¹¹

$$u(c_t | c_t^q) = \log c_t^q - \frac{(c_t / c_t^q)^{1-\rho} - 1}{\rho - 1},$$

¹⁰ To fully match the model, this state variable needs to reflect only recent experience: a formerly wealthy wine lover who has been drinking *vin de table* for many years may take a while to reacquire both a discerning palate and up-to-date knowledge of the wine market.

¹¹ This functional form, while not implied by the basic assumptions of the dual-self approach, is a standard CES and is not chosen with our particular data in mind.

where $\rho \geq 1$ corresponds to the short-run self's relative risk aversion over immediate consumption. Because our goal is a tightly parameterized model that fits the experimental data, this functional form assumes that all of the nightclubs have the same coefficient ρ . In practice ρ could vary. This might be important in accounting for the preferences of the very rich and the very poor; we discuss this in more detail when we examine the robustness of our calibrated parameters.

With this specification $u(c_t | c_t) = \log(c_t)$, and

$$\frac{\partial u(c_t | c_t^q)}{\partial c_t^q} = \frac{1}{c_t^q} - \left(\frac{c_t^q}{c_t}\right)^{\rho-2} \frac{1}{c_t}.$$

As a consequence, the first order condition for maximizing $u(c_t | c_t^q)$ with respect to c_t^q implies $c_t^q = c_t$, and the second order condition is

$$\left. \frac{\partial^2 u(c_t | c_t^q)}{\partial c_t^{q2}} \right|_{c_t=c_t^q} = -\frac{1}{c_t^2} - (\rho - 2) \frac{1}{c_t^2} = \frac{1}{c_t^2} (1 - \rho),$$

which is negative when $\rho > 1$.

Durable Consumption: The next step is to specify the agent's preferences for durable versus non-durable consumption. Our goal here is simply to account for the fact that durable consumption exists, and not to explain it, so we adopt a simple Cobb-Douglas-like specification $\tau u(c_t | c_t^q) + (1 - \tau) \log c_t^d$; this will lead to a constant share τ of spending on durables. Durable consumption c_t^d can only be adjusted slowly, and seems unlikely to respond at all to the sorts of income shocks received in the lab experiments we study. For this reason we simplify the model by assuming that the time path of c_t^d is chosen for once and for all in the initial time period.¹²

Cost of Self-Control: In our calibrations of the model, we assume that the cost function is quadratic: $g(v_t) = \gamma v_t + (1/2)\Gamma v_t^2$, so we will maintain that assumption throughout this paper. Thus the marginal cost of self-control at time t is $\gamma + \Gamma v_t$; we continue to denote this as γ_t .

¹² Similarly, we abstract from labor supply, precautionary savings motives, and so forth. However we explicitly introduce durable consumption so that when we calibrate pocket cash the short-run self does not perceive that the rent check and similar expenses are available for short-term amusement.

Long-run Self: The objective function of the long-run self is

$$U_{RF} = E \sum_{t=1}^{\infty} \delta^{t-1} [\tau (u(c_t | c_t^q) - g(u(x_t | c_t^q) - u(c_t | c_t^q))) + (1 - \tau) \log c_t^d]$$

which is to be maximized with respect to $c_t \geq 0, c_t^q \geq 0, c_t^d \geq 0, x_t \geq 0$ subject to w_1 given, $w_{t+1} = R(s_t + x_t - c_t)$, $s_t + x_t + c_t^d \leq w_t$ and $w_t \geq 0$. Notice that this is a simple optimization problem with no uncertainty and perfect foresight.

Solution in the Deterministic Case: This is essentially as before, except that we now choose the venue $c_t^q = c_t = x_t$ as well as the pocket cash equal to planned consumption. This means that the utility of the short-run self is $u(x_t | c_t^q) = \log x_t$, and as there is no self-control cost, this boils down to maximizing

$$\sum_{t=1}^{\infty} \delta^{t-1} [\tau \log x_t + (1 - \tau) \log c_t^d]$$

subject to the budget constraint $w_{t+1} = R(w_t - x_t - c_t^d)$. The solution is easily computed¹³ to be $x_t = (1 - \delta)\tau w_t$, $c_t^d = (1 - \tau)(1 - \delta)w_t$, which implies that $w_{t+1} = R(\delta w_t)$, as in the case without durable consumption. A calculation in the appendix shows that the corresponding present value utility of the long-run self is

$$U_1(w_1) = \frac{\log(w_1)}{1 - \delta} + K,$$

where

$$K = \frac{1}{1 - \delta} \left[\log(1 - \delta) + \frac{\delta \log(R\delta)}{(1 - \delta)} + \tau \log \tau + (1 - \tau) \log(1 - \tau) \right].$$

These together give the solution of the simple deterministic budget problem.

Because the time path of durable consumption c_t^d is chosen once and for all in period 1 as a function of initial wealth w_1 , the agent's problem in period 2 is to choose $\{x_t\}$ to maximize

¹³ The derivation is standard; an explicit computation in the case where $\tau = 1$ is in Fudenberg and Levine [2006]. Note that equation (1) of that paper contains a typographical error: in place of $(1 + \gamma)(\log(1 - a) + \log y_0)$ it should read $(1 + \gamma)\log(1 - a) + \log(y_0)$.

$$\sum_{t=2}^{\infty} \delta^{t-2} [\tau \log x_t + (1 - \tau) \log c_t^d]$$

subject to the wealth constraints $w_{t+1} = R(w_t - x_t - c_t^d)$ and $w_t \geq 0$.

We can treat the committed durable consumption as if it were all paid for in advance, so that the wealth available for non-durable consumption beginning in period 2 is $w_2 - (1 - \tau)R\delta w_1$.¹⁴ In the Appendix we use this to compute the maximized present value of utility

$$U_2(w_2 | w_1) = \frac{\tau \log(w_2 - (1 - \tau)R\delta w_1)}{1 - \delta} + \frac{(1 - \tau) \log((1 - \tau)R\delta w_1)}{1 - \delta} + K',$$

where

$$K' = \frac{1}{1 - \delta} \left[\log(1 - \delta) + \frac{\delta \log(R\delta)}{(1 - \delta)} \right].$$

In the Appendix where we prove Theorem 1, we show that it holds also in this augmented model, where now the relevant first order condition determining $\hat{c}_1^j(\gamma_1)(z_1^j)$ is

$$(c_1^j)^\rho = (c_1^q)^{\rho-1} \frac{(1 - \delta)(1 + \gamma_1)}{\delta} (\tau w_1 + z_1^j - c_1^j),$$

From the first order condition we may compute the cutoff

$$\hat{z}_1 = (c_1^q)^{\frac{\rho-1}{\rho}} \left[\frac{1 - \delta}{\delta} (1 + \gamma_1) [\tau w_1 - x_1] \right]^{1/\rho} - x_1.$$

7. Basic Calibration

The first step in our calibration of the model is to pin down as many parameters as possible using estimates from external sources of data. We will subsequently use data

¹⁴ Along the perfect-foresight path of the optimal plan, the wealth available for non-durable consumption is strictly positive. As we are now considering the optimal continuation for arbitrary values of period-two wealth w_2 , we must assume that the departure from the perfect-foresight path is small enough that this remains true. Of course it is possible for a consumer to lose so much wealth as to be unable to fulfill plans for future durable consumption, in which case they might become homeless, declare bankruptcy, and so forth. These events are far outside the range of any data we analyze here; the model is simply intended to be an idealization of a situation where these types of changes have a substantial fixed cost associated with them.

from laboratory experiments to calibrate risk aversion parameters and to determine the cost of self-control.

To measure the subjective interest rate r we ordinarily think of taking the difference between the real rate of return and the growth rate of per capita consumption. However, we must contend with the equity premium puzzle. From Shiller [1989], we see that over a more than 100 year period the average growth rate of per capita consumption has been 1.8%, the average real rate of returns on bonds 1.9%, and the real rate of return on equity 7.5%. Fortunately if the consumption lock-in once a nightclub is chosen lasts for six quarters,¹⁵ the problem of allocating a portfolio between stocks and bonds is essentially the same as that studied by Gabaix and Laibson [2001], which is a simplified version of Grossman and Laroque [1990].¹⁶ Their calibrations support a subjective interest rate of 1%. This rate, and any rate in the range 1-7%, can explain the data on the Allais paradox.¹⁷ To explain the Benjamin, Brown, and Shapiro data on Chilean high school students requires a higher interest rate of 7%, at the high end of what can be supported by Shiller's data.

From the Department of Commerce Bureau of Economic Analysis, real per capita disposable personal income in December 2005 was \$27,640. To consider a range of income classes, we will use three levels of income \$14,000, \$28,000, and \$56,000. To infer consumption from the data we do not use current savings rates, as these are badly mis-measured due to the exclusion of capital gains from the national income accounts. We instead use the historical long-term savings rate of 8% (see FSRB [2002]) measured when capital gains were not so important. This enables us to determine wealth and consumption from income.

We estimate wealth as annual consumption divided by our estimate of the subjective interest rate r : $w_1 = 0.92y_1 / r$, where y_1 denotes steady state income. In determining pocket cash, we need to take account of consumption c_t^d that is not subject

¹⁵ We have implicitly assumed it lasts for only a day, but the length of lock-in plays no role in the analysis, no result or calculation changes if the lock-in is six quarters.

¹⁶ They assume that once the nightclub is chosen, no other level of consumption is possible. We allow deviations from the nightclub level of consumption – but with very sharp curvature, so in practice consumers are “nearly locked in” to their choice of nightclub. Chetty and Szeidl [2006] show that these models of sticky consumption lead to the same observational results as the habit formation models used by Constantinides [1990] and Boldrin, Christiano and Fisher [2001].

¹⁷ Subjective interest rates outside this range are also consistent with the Allais data; we did not explore this as we focus on the range that has some prior support from macroeconomic data.

to temptation: housing, consumer durables, and medical expenses. At the nightclub, the rent or mortgage was already paid for at the bank, and it is not generally feasible to sell one's car or refrigerator to pay for one's impulsive consumption. As noted by Grossman and Laroque [1990], such consumption commitments increase risk aversion for cash gambles.¹⁸ For consumption data, we use the National Income and Product Accounts from the fourth quarter of 2005. In billions of current dollars, personal consumption expenditure was \$8,927.8. Of this \$1,019.6 was spent on durables, \$1,326.6 on housing, and \$1,534.0 on medical care, which are the non-tempting categories. This means that the share of income subject to temptation is $\tau = 0.57$. This number may be too large, as there are additional categories of consumption such as food, transportation, insurance, child-care, communication and health club dues that also may not be subject to temptation. For that reason we will also consider a robustness check with a much smaller value of τ .

Finally, we must determine the time horizon Δ of the short-run self. This is hard to pin down accurately, in part because it seems to vary both within and across subjects, but the most plausible period seems to be about a day. For the purposes of robustness we checked that none of our results are sensitive to assuming a time horizon of a week: details can be found in the earlier working paper version available on line.

Putting together all these cases, we estimate pocket cash to be

$$x_1 = 0.57 \times 0.92 \times y_1 / 365 = .00144 \times y_1.$$

Table 4 – Calibrated Parameter Summary

Percent interest r		$y_1 = 14K$		$y_1 = 28K$		$y_1 = 56K$	
annual	Daily	w_1	x_1	w_1	x_1	w_1	x_1
1	.003	1.3M	20	2.6M	40	5.2M	80
3	.008	.43M		.86M		1.7M	
5	.014	.30M		.61M		1.2M	

¹⁸ Chetty and Szeidl [2006] extend Grossman and Laroque to allow for varying sizes of gambles and costly revision of the commitment consumption. Postelwaite, Samuelson and Silberman [2006] investigate the implications of consumption commitments for optimal incentive contracts.

To determine a reasonable range of self-control costs, we need to find how the marginal propensity to consume “tempting” goods changes with unanticipated income. The easiest way to parameterize this is with the “self-control threshold,” which is the level of consumption at which self-control kicks in. The consumption cutoff corresponding to \hat{z}_1 is given by

$$\begin{aligned}\hat{c}_1 &\equiv x_1 + \hat{z}_1 = (x_1)^{\frac{\rho-1}{\rho}} \left[\frac{\tau(1-\delta)}{\delta} (1+\gamma_1) [w_2] \right]^{1/\rho} \\ &\approx x_1 (1+\gamma_1)^{1/\rho}\end{aligned}$$

where we use the facts that $w_2 \approx w_1 = x_1 / \tau(1-\delta)$, and that $\delta \approx 1$. Define $\mu_1(\gamma_1) = (1+\gamma_1)^{1/\rho} \approx \hat{c}_1 / x_1$. Because the numerical value of γ_1 is hard to interpret, we will report $\mu_1(\gamma_1)$ rather than γ_1 .

We can also relate μ_1 to consumption data. Abdel-Ghany et al [1983] examined the marginal propensity to consume semi- and non-durables out of windfalls in 1972-3 CES data.¹⁹ In the CES, the relevant category is defined as “inheritances and occasional large gifts of money from persons outside the family...and net receipts from the settlement of fire and accident policies,” which they argue are unanticipated. For windfalls that are less than 10% of total income, they find a marginal propensity to consume out of income of 0.94. For windfalls that are more than 10% of total income they find a marginal propensity to consume out of income of 0.02. Since the reason for the 10% cutoff is not clear from the paper, we will view 10% as a general indication of the cutoff.²⁰ As we have figured the ratio of income to pocket cash to be $y_1 / x_1 = 696$, the value of μ_1 corresponding to 10% of annual income is 69.6.

8. Small Stakes Risk Aversion

We begin by calibrating our model to data on small-stakes risk aversion. The central issue here, as emphasized by Rabin [2000] is that the small-stakes risk aversion

¹⁹ The Imbens, Rubin and Sacerdote [2001] study of consumption response to unanticipated lottery winnings shows that big winners earn less after they win, which is useful for evaluating the impact of winnings on labor supply. Their data is hard to use for assessing μ_1 , because lottery winnings are paid as an annuity and are not lump sum, so that winning reduces the need to hold other financial assets. It also appears as though the lottery winners are drawn from a different pool than the non-winners since winners earn a lot less than non-winners before the lottery.

²⁰ Landsburger [1966] with both CES data and with data on reparation payments by Germany to Israeli citizens reaches much the same conclusion.

observed in experiments implies implausibly large risk aversion for large gambles.²¹ Following Rabin, consider the choice between $(.5 : -100, .5 : 105)$ (“option A”) and 0 (“option B”); we expect that as Rabin predicts many people will choose B. Since the combination of pocket cash and the maximum winning is well below our estimates of \hat{c}_1 , all income is spent, and the consumer simply behaves as a risk-averse individual with wealth equal to pocket cash. Let us treat pocket cash as an unknown for the moment, and specialize our model to $\rho = 1$, so that $u(c, c^q) = \log(c)$ and the choice of nightclub is irrelevant. The consumer will reject the gamble if

$$.5 \log(x_1 - 100) + .5 \log(x_1 + 105) < \log(x_1),$$

and since such an x_1 exists, the dual-self model with logarithmic preferences can explain why the gamble is rejected. Moreover, calculation shows that the gamble will be rejected if pocket cash is less than \$2100, so the logarithmic model yields a sensible result here.²²

The problem with this analysis is that the gamble $(.5 : -100, .5 : 105)$ has comparatively large stakes. Laboratory evidence shows that subjects will reject considerably smaller gambles, which is harder to explain with short-run logarithmic preferences. We use data from Holt and Laury [2002], who did a careful laboratory study of risk aversion. Their subjects were given a list of ten choices between two lotteries S and L . The specific lotteries are shown below, where the first four columns show the probabilities of the rewards, and the first four rows, which are irrelevant to our analysis, are omitted.

²¹ Rabin thus expands on an earlier observation of Samuelson [1963].

²² Note that this theory predicts that if payoffs are delayed sufficiently, risk aversion will be much lower. Experiments reported in Barberis, Huang and Thaler [2003] suggest that there is appreciable risk aversion for gambles where the resolution of the uncertainty is delayed as well as the payoffs themselves. However, delayed gambles are subject to exactly the same self-control problem as regular ones, so this is consistent with our theory. In fact the number of subjects accepting the risky choice in the delayed gamble was in fact considerably higher than the non-delayed gamble, rising from 10% to 22%.

Table 5 – Laboratory Preferences Towards Risk

Option S		Option L		Fraction of Subjects Choosing S ²³			
\$2.00	\$1.60	\$3.85	\$0.10	1X	20X	50X	90X
Number of observations =>				(187)	(150)	(19)	(18)
0.5	0.5	0.5	0.5	.70(.03) 4.2	.85(.03)	1.0(0.0)	.90(.07)
0.6	0.4	0.6	0.4	.45(.04) 12	.65(.04)	.85(.08)	.85(.09)
0.7	0.3	0.7	0.3	.20(.03) 21	.40(.04)	.60(.12)	.65(.12)
0.8	0.2	0.8	0.2	.05(.02) 32	.20(.03)	.25(0.1)	.45(.12)
0.9	0.1	0.9	0.1	.02(.01) 50	.05(.02)	.15(.08)	.40(.12)
1.0	0.0	1.0	0.0	.00(0.0)	.00(0.0)	.00(0.0)	.00(0.0)

Bold numbers are coefficients of relative risk aversion giving indifference when

$$x_1 = 40$$

Initially subjects were told that one of the ten rows would be picked at random and they would be paid the amount shown. Then they were given the option of renouncing their payment and participating in a high stakes lottery, for either 20X, 50X or 90X of the original stakes, depending on the treatment. The high-stakes lottery was otherwise the same as the original: a choice was made for each of the ten rows, and one picked at random for the actual payment. Everyone in fact renounced their winnings from the first round to participate in the second. The choices made by subjects are shown in the table 5.

Assuming that the amounts of the lotteries lie below the threshold \hat{z}_1 , for each choice S, L and for given pocket cash x_1 we can compute the coefficient of relative risk aversion ρ that leads to indifference between the two choices.

²³ Standard errors in parentheses.

$$E(x_1 + \tilde{z}_1^S)^{1-\rho} = E(x_1 + \tilde{z}_1^L)^{1-\rho}.$$

These numbers are reported in bold face also in Table 5 for the base level of pocket cash $x_1 = 40$.

We use Table 5 to infer the sample median and the 85th percentile of risk aversion. For example, we can infer from the 1X column that the median individual prefers S when the probability of the high payment is 0.5, but prefers L when the probability of the high payment is 0.6. Hence the sample median value of ρ must lie between 4.2, which would lead to indifference when the probability is 0.5, and 12, which would lead to indifference when the probability is 0.6. Table 6 gives the ranges for the median value of risk aversion as well as the 85th percentile based on the data in Table 5.

Table 5: Bounds for Quantiles of Coefficient of Relative Risk Aversion

Pocket Cash $x_1 = 40$				
stakes	50 th percentile risk aversion		85 th percentile risk aversion	
1X	4.2	12	21	32
20X	1.1	1.9	2.8	4.2
50X	1.2	1.8	2.6	2.6
90X		1.5	2.1	

The first thing to observe is that the ranges for the 20X, 50X and 90X treatments are generally consistent with each other. The 1X treatments exhibit considerably higher risk aversion. This cannot be due to sampling error: Taking a relatively broad confidence band of three standard deviations, we see that 70% of subjects choose S in the 50-50 1X cell, three standard deviations below the mean is still 61% of subjects, so that the median relative risk aversion should be at least the 4.2 reported for that cell. On the other hand, in the 20X treatment, we see 20% of subjects choosing S in the 80-20 20X cell, and here three standard deviations above the mean is only 29% of subjects, so that the median should be no more than the 2.8 reported for that cell, well below the 4.2 minimum estimate for the 1X treatment. We conclude that the CES functional form does not fit very small stakes gambles particularly well. This is well known from the empirical

prospect theory literature, but is not a major concern, as we are not going to study gambles in the lower 1X range.²⁴

Examining the 20X, 50X and 90X ranges, for the median, we see that we can pin down the coefficient of relative risk aversion in the sample between a low value of 1.2 (highlighted in green) and a high value of 1.5 (highlighted in purple). Taking a non-parametric approach to estimating the corresponding population parameter in this case is unproblematic since the sample median is a consistent estimator of the population median.²⁵

For the 85th percentile – people who are more risk aversion than 85% of the population – the sample data is contradictory, as in the 50X treatment the coefficient of relative risk aversion is pinned down to 2.6, while the range in the 20X treatment is 2.8 to 4.2. However, the 50X treatment is subject to substantial sampling error: the estimate the 15% of subjects choosing S in 90-10 50X cell has a standard error of .08, so the information in that cell has little value. Hence the non-parametric approach gives a range from 2.8 to 4.2 for the 85th percentile.

In Table 4 we highlighted (in yellow and turquoise respectively²⁶) the cell that correspond most closely to the median and 85th percentiles. To choose a specific value from the ranges in Table 5, we choose values that fit these yellow (turquoise) cells as closely as possible. We use the CES utility scaled by the pocket cash constant x_1

$$-x_1 \frac{(c_1/x_1)^{1-\rho} - 1}{\rho - 1}.$$

For any given coefficient of relative risk aversion and any cell in Table 4, we can compute a squared utility difference between S and L . For the bold face coefficients shown in Table 4, this number will be zero, but we cannot choose a single value to simultaneously make the utility difference in all the cells equal to zero. Instead we compute the value of ρ that minimizes the sum of these squared utility differences. This

²⁴ It is possible that the size of the choices might have been confounded with the order in which the choices were given. Harrison, Johnson, McInnes and Rutstrom [2005] find that corrected for order the impact of the size of the gamble is somewhat less than Holt and Laurie found, a point which Holt and Laurie [2005] concede is correct. The follow-on studies which focus on the order effects do not contain sufficient data for us to get the risk aversion estimates we need.

²⁵ We could use a parametric model by assuming an explicit functional form for the population distribution of risk aversion, but such an estimator would be susceptible to specification error.

²⁶ Yellow cells for the median; turquoise cells for the 85th percentile.

yields a median of 1.43, which lies in the range 1.2 to 1.5 that we have already identified, and an 85th percentile of 2.74, which lies slightly below the range from 2.8 to 4.2.²⁷ We carry out the same computation for our other candidates for pocket cash \$20 and \$80, with the results reported in Table 6.

Table 6 – Estimated Relative Risk Aversion

	Pocket Cash x_1		
	\$20	\$40	\$80
ρ median	1.17	1.43	1.90
ρ 85 th	2.02	2.74	4.15

Because the estimated ρ 's are greater than 1, these preferences are not logarithmic. Notice that the data does not let us separately identify pocket cash and risk aversion; various combinations of these two are observationally equivalent.

We also wish to consider values of τ smaller than 0.57. Since we do not allow less risk aversion than the logarithm, the smallest possible value of pocket cash is that which gives rise to $\rho = 1$. We calibrate this to be $x = 16.4$, which corresponds to a value of $\tau = 0.23$ with the median income level. We use this value in robustness checks below.

9. The Allais Paradox

We proceed next to examine the Allais paradox in the calibrated model. We assume that the choice in this (thought) experiment is completely unanticipated. In this case the solution is simple: there is no self-control problem at the bank, so the choice is $c_1^* = x_1$ and spend all the pocket cash in the nightclub of choice. Given this, the problem is purely logarithmic, so the solution is to choose $x_1 = (1 - \delta)w_1$.

In the Kahneman and Tversky [1979] version of the Allais Paradox the two options in the first scenario are L_1 given by (.01 : 0,.66 : 2400,.33 : 2500), while S_1 is

²⁷ But also in the direction of the 90-10 50X cell.

2400 for certain. Many people choose option S_1 . In scenario two the pair of choices are $L_2 = (.67 : 0, .33 : 2500)$ and $S_2 = (.66 : 0, .34 : 2400)$.²⁸ Here many people choose $L_2 = .$ Expected utility theory requires the same option L or S be chosen in both scenarios.

From our qualitative analysis, we know that to get a reversal the safe alternative S must be preferred by u and the risky alternative L must be preferred by V . The first question is whether this is true for the calibrated parameters. Since it is, we know we can find a level of self-control cost that will lead to a reversal. As we will see, we cannot easily pin down the curvature of the self-control function, beyond the fact that it cannot be linear. However, the marginal cost of self-control must lead to near indifference between S and L , and so is well pinned down by the data. The second question to be addressed is whether this marginal cost is plausible and consistent with other data. Finally we carry out several comparative static exercises.

To describe the procedure we will use for reporting calibrations concerning choices between pairs of gambles, let us examine in some detail the choice between L and S in the base case where the annual interest rate $r = 1\%$, annual income is \$28,000, wealth is \$860,000, so pocket cash and the chosen nightclub are $x_1 = c_1^* = 40$. Recall the cost of self-control $g(v_1) = \gamma v_1 + (1/2)\Gamma v_1^2$. Consider first the case $\Gamma = 0$ of linear cost of self-control. Here we have an expected utility model, so the optimal choice is independent of the scenario, and we can solve for the numerically unique value γ_1^* ($\mu_1(\gamma_1^*) = 9.60$) such that there is indifference between the two gambles L and S .²⁹ A numerical computation shows that $Eu(\tilde{c}_1^S(\gamma_1^*)) > Eu(\tilde{c}_1^L(\gamma_1^*))$, so that when γ_1 is chosen so that the long-run self is indifferent, the short-run self prefers the sure outcome S . On the other hand, when there is no cost of self-control, it is easy to compute that the long-run self prefers the risky outcome L . This establishes that there will be an Allais type reversal.

²⁸ These were thought experiments. We are unaware of data from Allais experiments with real payments of over \$2000, though other experiments with stakes of analogous shares of average per-capita yearly income have been conducted in poor countries. There is experimental data on the Allais paradox with real, but much smaller, stakes; we discuss this below. We thank Chew Soo Hong for drawing our attention to an error in our discussion of this version of the paradox in a previous draft.

²⁹ We omit the subscript on L and S since here the independence axiom is satisfied, so it does not matter which pair of choices we consider. Subsequently when we add some curvature the indifference will be broken, and, as we shall see, in opposite ways for the first and second pair of choices.

To carry out the quantitative analysis, we computed γ_1^* for each of our cases, then constructed values of γ, Γ with γ close to γ_1^* together with corresponding marginal costs $\gamma_1^j[k]$ of each choice j in each scenario k that lead to a reversal. ³⁰

Table 7 Explaining the Allais Paradox with $r = 1\%$

income	$x_1 = c_1^*$	ρ	$\mu_1(\gamma_1^*)$	γ	Γ	$\mu_1(\gamma_1^L[1])$	$\mu_1(\gamma_1^S[1])$	$\mu_1(\gamma_1^L[2])$	$\mu_1(\gamma_1^S[2])$
14000	20	1.17	16.65	24.88	2	17.05	17.02	16.36	16.34
14000	20	2.02	7.56	56.11	32	7.62	7.60	7.47	7.45
28000	40	1.43	8.49	19.42	2	8.50	8.49	8.31	8.30
28000	40	2.74	4.13	45.07	64	4.16	4.15	4.10	4.09
56000	80	1.90	4.57	16.09	4	4.58	4.58	4.50	4.50
56000	80	4.15	2.47	39.37	128	2.48	2.47	2.46	2.45

³⁰ The programs used for the computations were in Octave, a free equivalent of Matlab. They can be found at www.dklevine.com/papers/allais.zip.




Figure 2 - Allais Paradoxes

Notice that for each case there is a wide range of parameters that will generate paradoxes. The basic limitation is that if the curvature Γ is very large, then it will be impossible to generate values of γ_1^L, γ_1^S that are sufficiently close to γ_1^* to give paradoxes. This is shown in Figure 2, where the parameter values generating Allais paradoxes are computed for the base case of \$28,000 annual income and risk aversion $\rho = 1.43$. In the blue-shaded region with low costs of self-control, the long-run optimum L is the best choice in both scenarios. In the red-shaded region with high costs of self-control, the short-run optimum S , is the best choice in both scenarios. In the green-shaded region in between an Allais reversal occurs, as the optimal choice is S in the high temptation scenario and L in the low temptation scenario.

The comparative statics are driven by how the marginal cost of self-control must change to maintain indifference between the two gambles in the face of changes in the marginal utilities of current and future consumption. Recall that the short-run self prefers

the safe gamble S , while, starting in period 2 the long-run self prefers the risky gamble L . If the marginal utility of current consumption increases relative to future consumption, this will break the tie in favor of the earlier period, that is, the short-run self. However, lowering the marginal cost of self-control effectively lowers the weight on the short-run self's preferences, and restores the tie. The short version: more weight on the present implies the marginal cost of self-control must fall lower in order to maintain indifference.

The first comparative static we consider is to change the interest rate. If we increase r to 3% or 5%, which increases the weight on the present, the marginal cost of self-control must fall to accommodate the paradox. However, in the calibration when we change r we also change wealth correspondingly. Changing r from 1% to 5% raises the weight on the present period by a little more than a factor of 5, but lowers wealth by a factor of 5, and since second period value is logarithmic, raises the marginal utility of second period wealth by a factor of 5. The net effect is a very small increase in the weight on the present, and when we did the calculation, the values of $\mu_1(\gamma_1^*)$ change only in the third significant digit.

In contrast, raising risk aversion holding everything else fixed makes the gamble less attractive to the short-run self, increasing the temptation. This effectively increases the weight on the first period, so must lead to a reduced marginal cost of self-control, as happens in Table 7. Increasing income has a different effect: it has little effect on the decision problem, since that is formulated in relative terms. That means that the cutoff in dollars cannot change much, and so the cutoff relative to pocket cash, which has increased, must go down.

The values of the self-control parameter $\mu_1(\gamma_1^*)$ range from 2.47-16.65, which is considerably smaller than the 69.6 figure Abdel-Ghany et al [1983] found in CES data that we discussed above. However, the windfall income in the CES is considerably larger than these Allais gambles, so poses a greater temptation, and with increasing marginal cost of self-control should generate higher marginal self-control costs and thus even higher thresholds.

Notice that we are able to explain the Allais paradox with exactly the same risk aversion parameters that we used to explain the Rabin paradox. The theory here give a consistent explanation of both paradoxes, and it does so with a decision model that is

consistent with long-run savings behavior being logarithmic as in growth and macroeconomic models.

As a robustness check, we also examined what happens for the median income of \$28,000 when we use a much smaller value $\tau = 0.23$ with the corresponding risk aversion $\rho = 1$.

Table 8 Explaining the Allais Paradox with $\tau = 0.23$

income	$x_1 = c_1^*$	ρ	$\mu_1(\gamma_1^*)$	γ	Γ	$\mu_1(\gamma_1^L[1])$	$\mu_1(\gamma_1^S[1])$	$\mu_1(\gamma_1^L[2])$	$\mu_1(\gamma_1^S[2])$
28000	16.4	1.00	24.31	22.15	1	24.95	24.92	23.80	23.77

As can be see the results are very similar to the base case.

Now we examine whether these parameters are consistent with the Keren and Roelofsma [1995] data on hyperbolic discounting reported above. The non-linearity here is not sufficient to generate a reversal in the Keren and Roelofsma experiment: Computation shows that even with the 50% chance of a prize, an individual with any of the parameters in the table above strictly prefers to take the money now. The reason is that our estimates of the linear coefficient γ are too large to support a reversal. If we use the same value of the curvature coefficient Γ as before, but a lower intercept, then a reversal is generated.³¹ However, this lower value of the linear coefficient cannot explain the Allais paradox, because it reduces the temptation of the Allais gambles so much that the riskier gamble L is preferred in both scenarios. It might be possible to accommodate both the Allais choices and Keren and Roelofsma data by departing from the quadratic specification of control costs, but we have not explored this possibility.

Original Allais Paradox: The original Allais paradox involved substantially higher stakes, so it would be difficult to implement other than as a thought experiment: option L_1 was $(.01 : 0, .89 : 1,000,000, .1 : 5,000,000)$ and S_1 was 1,000,000 for certain,; the second scenario was $L_2 = (.90 : 0, .10 : 5,000,000)$ and $S_2 = (.89 : 0, .11 : 1,000,000)$.

³¹ For example, with low income, high risk aversion, and an annual subjective interest rate of 3%, if the intercept is taken to be 4.4 rather than 26.4, and we use the estimated curvature of 23.9, a reversal is generated for the Keren and Roelofsma [1995] data.

Here also the paradoxical choices were S_1 and L_2 . In our base case of median income the results for the original Allais paradox are reported in Table 9.

Table 9: Explaining the Original Allais Paradox with $r = 1\%$

income	$x_1 = c_1^*$	ρ	$\mu_1(\gamma_1^*)$	γ	Γ	$\mu_1(\gamma_1^L[1])$	$\mu_1(\gamma_1^S[1])$	$\mu_1(\gamma_1^L[2])$	$\mu_1(\gamma_1^S[2])$
28000	40	1.43	4336	155148	262144	4408	4308	4301	4198
28000	40	2.74	128	549929	9961472	130	125	130	125

Notice that the values of $\mu_1(\gamma_1^*)$ are considerably larger here than they are for the Kahneman-Tversky version of the paradox. This is as it should be: $\mu_1(\gamma_1^*)$ is endogenous and determined by temptation. The original Allais paradox involves larger stakes and thus larger temptations than the Kahneman-Tversky version, so the theory predicts that the marginal cost of self-control should be larger. The value corresponding to low risk aversion, however, is implausible. It implies that the cutoff in dollars is about \$400,000, meaning that if the outcome is favorable (winning either \$1,000,000 or \$5,000,000) the long-run self intends to allow the short-run self to spend this amount on the first day. The value corresponding to high risk aversion is still large but more sensible: in case of success the long-run self will allow the short-run self to spend \$5200 immediately. Since our model was calibrated on data with real payoffs, and that original paradox involves very large amounts that subjects may find difficult to evaluate, the discrepancy does not seem like a major concern.

10. Cognitive Load

The theory predicts that increasing cognitive load should increase the marginal cost of self-control and lead to reversals similar to those in the Allais paradox. Relatively few experiments have been conducted on the effect of cognitive load on decisions involving risk. One recent one is an experiment conducted with Chilean high school juniors by Benjamin, Brown and Shapiro [2006]. We analyze their data to show that their subjects have Allais-like reversals brought about by cognitive load as predicted by the theory and that quantitatively this occurs in our calibrated model.

In the experiment students made choices about uncertain outcomes both under normal circumstances and under the cognitive load of having to remember a seven-digit number while responding. In scenario 1 the choice was between a 50-50 gamble between 650 pesos and nothing versus a sure option of 250 pesos. In scenario 2 the sure option was replaced by a 50-50 gamble between 300 and 200 pesos.³² Table 10 below summarizes the fraction of the population taking the risky choice, with the number in parentheses following the treatment indicating the number of subjects.

Table 10 - Students Taking the High Risk Option

650/0 versus 250		650/0 versus 300/200	
No load (13)	Cognitive Load (21)	No Load (15)	Cognitive Load (22)
70%	24%	73%	68%

These were real, and not hypothetical choices, the subjects were paid in cash at the end of the session. To provide some reference for these numbers, 1 \$US= 625 pesos and the subjects average weekly allowance was around 10,000 pesos from which they had to buy themselves lunch twice a week.³³

The key fact in the table is that introducing cognitive load when the alternative is safe induces many subjects to switch to the safe alternative, while there is no such reversal when the “safe” alternative is the 300/200 gamble. This is as the theory predicts. If the short-run self prefers the safe alternative to the risky one we should see the first reversal. However, the 300/200 gamble is less tempting than the sure alternative of 250, so a cognitive load that will lead to a reversal in the first scenario need not do so in the second.

To calibrate the model, we take pocket cash to be the average weekly allowance of 10,000 pesos divided by 7 that amount in the daily case, or about \$2.29. We then work out wealth and income indirectly using the utility-function parameters that we calibrated

³² We thank the authors for providing us with this data. There is data on a risky alternative involving four other size prizes that are not relevant for our purposes. There is one anomaly in the data that we cannot explain: the fraction of people choosing the risky option against the sure alternative under cognitive load actually decreases as the size of the prize is increased. This may be due to sampling error.

³³ Many of them buy lunch at McDonald’s for 2000 pesos twice a week, leaving an apparent disposable income of 6000 pesos per week.

in the Allais experiments.³⁴ To explain a preference reversal, the parameters must lead the two choices to have sufficiently similar levels of utility so that self-control matters. Within the range of parameters in our calibration, the only set of parameters for which this is true is when the annual interest rate is 7% and risk aversion is at the lower median level. For these parameters, we can calculate the values of $\mu_1(\gamma_1^*)$ that leads to indifference in the first and second scenario respectively shown in Table 11.

Table 11 - Parameters for Indifference for the Chilean Gambles

r	income	w_1	$x_1 = c_1^*$	ρ	$\mu_1(\gamma_1^*)$ 1	$\mu_1(\gamma_1^*)$ 2
7%	1.6K	21K	2.29	1.17	21.219	21.222

Note that the values of $\mu_1(\gamma_1^*)$ of 21.219-21.222 needed to create indifference for the Chilean gambles are close to the value $\mu_1(\gamma_1^*) = 19.2$ from the Allais paradox for the 5% calibration. This makes sense, because the temptations are of the same order of magnitude.

In both scenarios, the risky option has the greater temptation, meaning that it will be chosen only for low marginal cost of self-control or equivalently, low values of γ_1^* . The risky option, however, is preferred in the absence of cost of self-control. Recall that in our model the marginal cost of self-control is $\gamma + \Gamma(d_1 + \bar{u}_1 - u_1)$ where d measures the cognitive load. Suppose that $\mu_1(\gamma_1) < 21.219$ and that Γ is not too large. Then when cognitive load $d_1 = 0$, marginal cost of self-control is low enough in both scenarios that the risky alternative will be chosen. On the other hand, when cognitive load is high so $d_1 = \bar{d}_1 > 0$, for an appropriate value of \bar{d}_1 , there will be a greater marginal cost of self-control $21.219 \leq \mu_1(\gamma_1) \leq 21.222$. That means that in scenario 1 the marginal cost of self-control is high enough that the safe alternative will be chosen, while in scenario 2 the marginal cost of self-control is low enough so that the risky alternative will continue to be chosen.

Note that the cognitive load calibration is more sensitive to the interest rate than the Allais calibration, as with cognitive load we require that r be at least 5%. In both cases, increasing the interest rate slightly increases the weight on the present relative to

³⁴ It is unclear that we should use the same value of τ but the results are not terribly sensitive to this.

the future, due to the offsetting effect of changing wealth in the calibration. The intuition for this is that in the Allais case it is the levels of the marginal cost of self-control that matters, while in the cognitive load case what matters is the difference between two different marginal costs of self-control. The latter is much smaller than the absolute level, and so smallish increases in the level can result in largish increases (proportionally) in the difference.

In both cases, increasing the interest rate increases by a small amount (due to the offsetting effect of changing wealth in the calibration) the weight on the present relative to the future. In both cases this has the effect of slightly reducing the cost of self-control that leads to indifference. In the cognitive load calibration, there is less temptation in scenario 2 than scenario 1, meaning that for indifference $\mu_1(\gamma_1^*)$ is larger in scenario 2 than scenario 1, as required to explain the data. However, for $r = 1\%, 3\%$, the weight on the first period is so small that the algorithm is unable to find indifference. As we increase the weight on the early period, the amount by which we must adjust self-control to maintain indifference for a given drop in temptation increases. At $r = 5\%$ the computer can find it, and the gap expands considerably as we increase the interest rate further. For interest rates higher than 5% – not implausible for high school students – we find that the range expands farther, making it more likely that cognitive load could push the marginal cost of self-control into the intermediate range needed to explain the data.

11. Limitations of the Theory

The model predicts that reversals will be attenuated by delay, and we do find this, for example, in the Keren and Roelofsma data. However, the calibrated model is excessively stark: it predicts that the percentage of reversals should fall to zero as soon as there is a 24 hour delay. This is a more general problem with the model (shared with quasi-hyperbolic discounting): There is a great deal of data showing that the effect of delay is gradual and not abrupt. (See for example, Green and Myerson [1996].)

One possibility is that there may be population heterogeneity in the horizon of the short-run self. But in the Green and Myerson data we see gradual attenuation even for a single individual. To model gradual delay at the individual level in this framework it is necessary that short-run selves live more than one period. There are two possibilities here. One is that there are overlapping generations of short-run selves. Such a framework

is not altogether attractive since it introduces strategic interactions between the short-run selves akin to those in hyperbolic discounting theory and loses the basic simplicity of a single time-consistent optimization problem. An alternative is to assume that short-run selves live random lengths of time – for example face a fixed hazard rate of dying and being replaced by another short-run self each period. Whether such a model could explain quantitatively the gradual effects of delay remain for future research.

Next we turn to the quantitative predictions of the model. Although the base model qualitatively predicts the common ratio paradox, it cannot explain quantitatively the Baucells and Heukamp data, at least not for costs of self-control comparable to those used for the Allais paradox. The reason simply is that the payoffs are all below the threshold $\mu_1(\gamma_1)$ (roughly 1% of annual income) used to explain the Allais paradox. This means that the decision-maker simply uses the preferences of the short-run self, and so behaves as an expected utility maximizer with wealth equal to pocket cash.

Just as this stark model cannot explain small stakes common ratio paradoxes, it cannot explain small stakes Allais paradoxes. Whether these paradoxes exist is somewhat controversial. Battalio, Kagel and Jiranyakul [1990] found that subjects did exhibit the Allais paradox even for very small stakes in the range \$0.12 to \$18.00, and that they less frequently even made the reverse Allais choices. However, indifference or near indifference may be a key factor in the reported results. In set 1 and set 2 of Battalio et al the two lotteries have exactly the same expected value, the difference between the large and small prize is at most \$8.00, and there was only one chance in fifteen that the decision would actually be implemented. So it is easy to imagine that subjects did not invest too much time and effort into these decisions.

By way of contrast Harrison [1994] found that with various small stakes the Allais paradox was sensitive to using real rather than hypothetical payoffs, and found in the real payoff case only 15% of the population exhibited the paradox. Although Camerer³⁵ pointed out the drop from 35% when payoffs were hypothetical was not statistically significant, a follow study by Burke, Carter, Gominiak and Ohl [1996] found a statistically significant drop from 36% to 8%. Conlisk [1989] also finds little evidence of an Allais paradox when the stakes are small. He examines payoffs on the order of \$10,

³⁵ Personal communication to the authors.

much less than our threshold values of $\mu_1(\gamma_1)$ of roughly 1% of annual income. These studies suggest that when played for small real stakes the Allais paradox is rare, as the stark version of the theory predicts.

Why is it that our model is not consistent with small-stakes reversals? The key fact is that when the stakes are small enough the individual is at the corner where no self-control is used, and so behaves as a purely short-run expected utility maximizer. Our simple and stark model rules out any incentive for self-control when stakes are small, but this is not entirely plausible, nor is it difficult to adapt the model so that some self-control is used. For example, if we assume that there is a small probability the bank will be closed each period, the optimum will be to use a little self-control to save a little bit as a precaution against bank closure. In this case we will get Allais type reversals even with small stakes.

12. Alternative Theories

The existing model most widely used to explain a variety of paradoxes, including the Allais paradox, is prospect theory,³⁶ which involves an endogenous reference point that is not explained within the theory.³⁷ In a sense, the dual-self theory here is similar to prospect theory in that it has a reference point, although in our theory the reference point is a particular value, pocket cash. The key aspect of pocket cash is that it is not arbitrary, but is endogenous and depends in a specific way on the underlying preference parameters of the individual. The theories are also quite different in a number of respects. Prospect theory makes relatively *ad hoc* departures from the axioms of expected utility, while our departure is explained by underlying self-control costs. Our theory violates the independence of irrelevant alternatives, with choices dependent on the menu from which choices are made, while prospect theory satisfies independence of irrelevant alternatives. Our theory can address issues such as the role of cognitive load and explains intertemporal paradoxes such as the hyperbolic discounting phenomenon and the Rabin

³⁶ There are many other alternative explanations of specific paradoxes, for example regret theory but these are less comprehensive than either this theory and prospect theory. There are also paradoxes that are hard to explain in any theory. For example Sydnor [2006] shows that a typical homeowner pays \$100 to reduce his deductible from \$1000 to \$500 with a claim rate of less than 5%. Rather than revealing something about risk preference, this may be the result of the homeowner's not knowing the probability of a loss.

³⁷ See Kozegi and Rabin [2006] for one way to make the reference point endogenous, and Gul and Pesendorfer [2007] for a critique.

paradox about which prospect theory is silent. Finally, a primary goal of our theory is to have a self-contained theory of intertemporal decision making; by way of contrast, it is not transparent how to embed prospect theory into an intertemporal model.³⁸

In the other direction, prospect theory allows for individuals who are simultaneously risk averse in the gain domain and risk loving over losses. This is done in part through the use of different value functions in the gain and loss domains, and in part through its use of a probability weighting function, which can allow individuals to overweight rare events.³⁹ Most work on prospect theory has estimated a representative-agent model; Bruhin, Fehr-Duda, and Epper [2007] refined this approach by classifying individuals as expected utility maximizing or as cumulative prospect theory types,⁴⁰ and find that most individuals are prospect theory types. It is interesting to note that given the functional forms they estimate, individuals with expected utility preferences are assumed to be risk averse throughout the gains domain, while in their data individuals are risk loving for small probabilities of winning, while for higher probability of success they are risk averse.

The stark model presented here cannot explain risk-loving choices over certain ranges. On the other hand, risk seeking can be explained within the expected utility paradigm by means of a Savage-style S-shaped utility function that is risk loving for small increases in income and risk averse for larger increases,⁴¹ and we can incorporate such S-shape utility functions into our theory. The only requirement on the short-run utility function (as a function of choice of nightclub) is that its upper envelope be the long-run logarithmic utility function; this forces the short-run utility to be at least as risk averse as the logarithm at $c = c^q$. Away from that optimal choice of nightclub it allows a wide variety of possibilities. This includes S shaped curves, and risk aversion that varies with wealth, or that exhibits heterogeneity independent of long-run risk aversion.

³⁸ Kozegi and Rabin [2007] develop but do not calibrate a dynamic model of reference dependent choice.

³⁹ See Prelec [1998] for an axiomatic characterization of several probability weighting functions, and a discussion of their properties and implications.

⁴⁰ Their estimation procedure tests for and rejects the presence of additional types.

⁴¹ Notice that it is possible to embed such short-run player preferences in our model although we have focused on the risk averse case. Indeed, such preferences are consistent even with long-run risk aversion: the envelope of S-shaped short-term utility functions can be concave provided that there is a kink between gains and losses, with strictly higher marginal utility in the loss domain. There is evidence that this is the case.

While S-shaped utility can explain risk seeking for small chances of gain and risk aversion for larger chances, it does not explain the Allais paradox, while prospect theory can potentially do so. But it appears that the parameters needed to explain individuals who are simultaneously risk averse and risk loving cannot at the same time explain the Allais paradox. Neilson and Stowe [2002] conducted a systematic examination of the parameters needed to fit cumulative prospect theory (with rank-dependent probabilities) to various empirical facts, and concluded that

parameterizations based on experimental results tend to be too extreme in their implications. The preference function estimated by Tversky and Kahneman (1992) implies an acceptable amount of risk seeking over unlikely gains and risk aversion over unlikely losses, but can accommodate neither the strongest choice patterns from Battalio, Kagel, and Jiranyakul (1990) nor the Allais paradox, and implies some rather large risk premia. The preference functions estimated by Camerer and Ho (1994) and Wu and Gonzalez (1996) imply virtually no risk seeking over unlikely gains and virtually no risk aversion over unlikely losses, so that individuals will purchase neither lottery tickets nor insurance.... We show that there are no parameter combinations that allow for both the desired gambling/insurance behavior and a series of choices made by a strong majority of subjects and reasonable risk premia. So, while the proposed functional forms might fit the experimental data well, they have poor out-of-sample performance.

The point is simply that prospect theory does not explain the data used to motivate it with a single consistent set of preferences. Notice, as we discussed above, that the stark version of the dual-self model presented here also cannot explain risk seeking over unlikely gains, nor does our CES specification fit well on very low stakes. Neilson and Stowe [2002] examined the original Allais paradox holding relative risk aversion constant, which as we have already noted is quite difficult because with expected utility individuals are not near indifferent in these large-stakes gambles if they have reasonable degrees of risk aversion. However, if we use the Bruhin, Fehr-Duda, and Epper [2007] estimates from the Zurich 03 gains- domain treatment, the prospect theory types have preferences give by

$$U = \sum_i \frac{.846p_i^{.414}}{.846p_i^{.414} + (1 - p_i)^{.414}} x_i^{1.056}$$

where p_i is the probability of winning the prize x_i .⁴² In the Kahnemann and Tversky version of the Allais paradox, L_1 is (.01 : 0,66 : 2400,33 : 2500), and S_1 is 2400 for certain. This gives $U(L_1) = 3874.58$ and $U(S_1) = 3711$. In other words, an individual with these preferences would prefer L_1 to S_1 and so would not exhibit an Allais paradox.

13. Conclusion

We have argued that a simple self-control model with quadratic cost of self-control and logarithmic preferences can account quantitatively for both the Rabin and Allais paradoxes. We have argued also that the same model can account for risky decision making of Chilean high school students faced with differing cognitive loads. Ranges of income from half to double the median income; subjective interest rates in the range of 1-7%; short-run risk aversion in the range from 1-4; and a self-control cost switchpoint $\mu_1(\gamma_1^*)$ in the range 15-30 cover all of the cases. Except for the Chilean data, these results are quite robust. The Chilean students' behavior, however, require high subjective interest rates of 7%.

We find it remarkable that the behavior of Chilean high school students can be explained with essentially the same parameters that explain the Allais paradox. This finding is not trivial, as there are possible observations that are not consistent with the theory. For example, cognitive load in the Chilean experiment could have caused subjects to switch in the reverse, "anti-Allais," direction, which our model could not explain. Second, there is enormous heterogeneity in the data; only a fraction of subject populations exhibit reversals, and the populations in the various experiments are very different, so there is no reason to believe that there is a single set of individual parameter values that will explain all of the data. However, while we have allowed the parameters

⁴² Bruhin, Fehr-Duda, and Epper [2007] specified a utility function only for two-outcome gambles, this seems the natural extension to the three or more outcomes demanded to explain the Allais paradox. Note also that this utility function has the highly unlikely global property that if we fix the probabilities of the outcomes and vary the size of the rewards it exhibits strict risk loving behavior.

to vary somewhat across experiments, it is important that all the parameters we use fall within a plausible range.

The main tension in explaining the experimental data using parameters calibrated to aggregate savings data is with respect to the degree of self-control. The model predicts that there should be a threshold level of unanticipated income, with marginal propensity to consume of 100% below the threshold and a very low marginal propensity above it. As we indicated, the permanent consumption data analyzed by Abdel-Ghany et al [1983] indicates that this may be true, and that the threshold is about 10% of annual income or 69.6 times pocket cash. We find, however, that to explain the data we consider, the threshold must be in the range of 4.04-22.3, which is considerably smaller than the threshold found in household consumption surveys. This is consistent with our model, as the windfalls in the CES data were larger than those in the experiments, and with more temptation and increasing marginal cost of self-control, it is optimal to allow the short-run self to spend more.

Our overall summary, then, is that the dual-self model explains choices over lotteries about as well as prospect theory, while explaining phenomena such as commitment and cognitive load that prospect theory cannot. Moreover, the dual-self model is a fully dynamic model of intertemporal choice that is consistent with both traditional models of savings (long-run logarithmic preferences) and with the equity premium puzzle.⁴³

In conclusion, there is no reason to think that the dual-self model has yet arrived at its best form, but its success in providing a unified explanation for a wide range of phenomena suggests that it should be viewed as a natural starting point for attempts to explain other sorts of departures from the predictions of the standard model of consumer choice. One possible next step would be try to more explicitly account for the evident heterogeneity of the population, and estimate distributions of self-control parameters as opposed to simply fitting the median or some other fractile as we have done here.

⁴³ The “behavioral life cycle model” of Shefrin and Thaler [1988] can also explain many qualitative features of observed savings behavior, and pocket cash in our model plays a role similar to that of “mental accounts” in theirs. The behavioral life cycle model takes the accounts as completely exogenous, and does not provide an explanation for preferences over lotteries. It does seem plausible to us that some forms of mental accounting do occur as a way of simplifying choice problems. In our view this ought to be derived from a model that combines the long-run/short-run foundations of the dual-self model with a model of short-run player cognition.

References

- Abdel-Ghany, M., G. E. Bivens, J. P. Keeler, and W. L. James, [1983], “Windfall Income and the Permanent Income Hypothesis: New Evidence,” *Journal of Consumer Affairs* 17: 262–276.
- Barberis, Huang and Thaler [2003], “Individual Preferences, Monetary Gambles, and the Equity Premium,” NBER WP 9997.
- Battalio, R. C., J. H. Kagel, and K. Jiranyakul [1990], “Testing Between Alternative Models of Choice and Uncertainty: Some Initial Results,” *Journal of Risk and Uncertainty*, 3: 25-90.
- Baucells, M., F. H. Heukamp and A. Villasis [2007], “Probability and Time Tradeoff: Experimental Evidence,” *IESE Business School Working Paper*.
- Benjamin, D., S. Brown, and J. Shapiro [2006], “Who is Behavioral: Cognitive Ability and Anomalous Preferences,” mimeo.
- Benhabib, J. and A. Bisin, [2005], “Modeling Internal Commitment Mechanisms and Self-Control: A Neuroeconomics Approach to Consumption-Saving Decisions.” *Games and Economic Behavior* 52:460-492.
- Bernheim, B. Douglas and A. Rangel [2004], “Addiction and Cue-Triggered Decision Processes.” *American Economic Review*, 94:1558-1590.
- Boldrin, M., L. Christiano and J. Fisher [2001], “Habit Persistence, Asset Returns, and the Business Cycle,” *American Economic Review* 91:149-166.
- Brocas, I. and J. Carrillo [2005], “The Brain as a Hierarchical Organization,” mimeo.
- Bruhin, A., H. Fehr-Duda, and T. Epper [2007], “Risk and Rationality: Uncovering Heterogeneity in Probability Distortion,” mimeo.
- Burke, M. S., J. R. Carter, R. D. Gominiak and, D. F. Ohl [2006], “An experimental note on the allais paradox and monetary incentives,” *Empirical Economics*.
- Chetty, R. and A. Szeidl [2007] ,“Consumption Commitments and Risk Preferences,” *Quarterly Journal of Economics* 122:831-877.
- Camerer, C. F. and T-H. Ho [1994] “Violations of the Betweenness Axiom and Nonlinearity in Probability” *Journal of Risk and Uncertainty* 8: 167-196.

- Conlisk, J. [1989], "Three Variants on the Allais Example," *American Economic Review*, 79: 393-407
- Constantinides, G. [1990], "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy* 98:519-543.
- Cox, J., V. Sadiraj, B. Vogt, and U. Dasgupta [2007], "Is There a Plausible Theory for Risky Decisions?" Georgia State University
- Dekel, E. [1986], An axiomatic characterization of preferences under uncertainty: Weakening the independence axiom, *Journal of Economic Theory*, 40: 304-318.
- Dekel, E. , Lipman, B., and Rustichini, Aldo. [2008], "Temptation-Driven Preferences." forthcoming in the *Review of Economic Studies*.
- Frederick, S. [2005], "Cognitive Ability and Decision Making," *Journal of Economic Perspectives*, 19:24-42.
- FRBSF [2002], "What's Behind the Low U.S. Personal Saving Rate?" *FSRB Economic Letter*, 2002-09; March 29, 2002.
- Fudenberg, D. and D.K. Levine [2006], "A Dual Self Model of Impulse Control," *American Economic Review* 96: 1449-1476.
- Gabaix, X. and D. Laibson [2001], "The 6D Bias and the Equity Premium Puzzle", *NBER Macroeconomics Annual*.
- Gailliot, M., R. F. Baumeister, C. N. DeWall, J. K. Maner, E. A. Plant, D. M. Tice, L. E. Brewer, [2008], "Self-Control Relies on Glucose as a Limited Energy Source: Willpower Is More Than a Metaphor," *Journal of Personality and Social Psychology*, forthcoming
- Green, I. And J. Myerson [1996], "Exponential Versus Hyperbolic Discounting of Delayed Outcomes: risk and Waiting Time," *American Zoologist*, **36**: 496-505.
- Grossman, S. and G. Laroque [1990], "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica* 58: 25-52
- Gul, F. and W. Pesendorfer [2001], "Temptation and Self-control." *Econometrica*, 69:1403-1436.
- Gul, F. and W. Pesendorfer [2005], "The Case for Mindless Economics," mimeo.
- Halevy, Y. [2008], "Strotz Meets Allais: Diminishing Impatience and the Certainty Effect," *American Economic Review*, **98** 1145-1162.

- Harrison, G. W. [1994], "Expected Utility Theory and the Experimentalists," *Empirical Economics*, 19: 223-243
- Harrison, G. W., E. Johnson, M. M. McInnes, E. E. Rutström [2005], "Risk Aversion and Incentive Effects: Comment," *American Economic Review* 95:897-901.
- Holt, C. and S. Laury [2002], "Risk Aversion and Incentive Effects," *American Economic Review* 92:1644-1655.
- Holt, C. and S. Laury [2005], "Risk Aversion and Incentive Effects: New Data Without Order Effects" *American Economic Review* 95:902-904.
- Imbens, G., D. Rubin, B. Sacerdote [2001], "Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players," *American Economic Review* 91:778-794.
- Kahneman, D. and A. Tversky [1979], "Prospect Theory: An Analysis of Decisions under Risk," *Econometrica* 47, 313-327.
- Keren, G. and P. Roelofsma [1995], "Immediacy and Certainty in Intertemporal Choice," *Organizational Behavior and Human Decision Making*, 63 297-297.
- Koszegi, B. and M. Rabin [2006], "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics* 121:1133-1165.
- Koszegi, B. and M. Rabin [2007], "Reference-Dependent Consumption Plans," mimeo.
- Laibson, D. [1997], "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112: 443-477.
- Landsberger, M. [1966], "Windfall Consumption: Comment", *American Economic Review*, 56: 534-540
- Loewenstein, G. and T. O'Donoghue [2004], "Animal Spirits: Affective and Deliberative Processes in Economic Behavior," mimeo.
- Muraven, M. and R. Baumeister [2000], "Self-Regulation and Depletion of Limited Resources: Does Self-Control Resemble a Muscle?" *Psychological Bulletin*, 126 247-259.
- Muraven, M. M. D. Tice and R. Baumeister [1998], "Self-Control as a Limited Resource." *Journal of Personality and Social Psychology*, 74 . 774-789.
- Neilson, W. and J. Stowe [2002], "A Further Examination of Cumulative Prospect Theory Parameterizations," *Journal of Risk and Uncertainty* 24:31-46.

- Ozdenoren, E., S. W. Salant and D. Silverman [2006], "Willpower and the Optimal Control of Visceral Urges" NBER W12278.
- Postelwaite, A., L. Samuelson, and D. Silberman [2004], "Consumption Commitments and Preferences for Risk," mimeo.
- Prelec, D. [1998], "The Probability Weighting Function," *Econometrica* 66:497-527.
- Rabin, M. [2000], "Risk Aversion and Expected utility Theory: A Calibration Theorem." *Econometrica*, 68:1281-1292.
- Samuelson, P. [1963], "Risk and Uncertainty: A Fallacy of Large Numbers," *Scientia* 98:108-113.
- Shefrin, H. and R. Thaler [1988], "The Behavioral Life-Cycle Hypothesis," *Economic Inquiry*, 26:609-643.
- Shiv, B. and A. Fedorikhin [1999], "Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making." *Journal of Consumer Research*, 26: 278-292.
- Shiller, R. [1989], *Market Volatility*, MIT Press: Cambridge.
- Sydnor, J. [2006], "Sweating the small stuff: The demand for low deductibles in homeowners insurance," mimeo, UC Berkeley.
- Thaler, R. H. [1980], "Towards a positive theory of consumer choice," *Journal of Economic Behavior and Organization*, 1:39-60
- Ward, A. and T. Mann, Traci [2000], "Don't Mind If I Do: Disinhibited Eating Under Cognitive Load." *Journal of Personality and Social Psychology*, 78: 753-763.
- Weber, B. and G. Chapman [2005], "The Combined Effects of Risk and Time on Choice: Does Uncertainty Eliminate the Immediacy Effect? Does Delay Eliminate the Certainty Effect?" *Organizational Behavior and Human Decision Processes*, **96**, 104-118

Appendix

The text considers four cases: the *general case* of nested utility and durable consumption, and two special cases: *nested CES/log*, the *base case* of non-nested utility without durable consumption, and the *base case specialized to logarithmic* utility. Here we prove the results stated in the text, with some generalizations.

Calculation of $V_1(w_1)$ and $V_2(w_2 | w_1)$ with nested CES/log: Recall from the text that

$$x_t = (1 - \delta)\tau w_t, \quad c_t^d = (1 - \tau)(1 - \delta)w_t, \quad w_{t+1} = R(w_t - x_t - c_t^d) \quad \text{so} \quad w_{t+1} = R(\delta w_t)$$

and thus $w_t = (R\delta)^{t-1}w_1$. We may then compute

$$\begin{aligned} V_1(w_1) &= \sum_{t=1}^{\infty} \delta^{t-1} [\tau \log x_t + (1 - \tau) \log c_t^d] \\ &= \sum_{t=1}^{\infty} \delta^{t-1} [\tau \log((1 - \delta)\tau w_t) + (1 - \tau) \log((1 - \tau)(1 - \delta)w_t)] \\ &= \sum_{t=1}^{\infty} \delta^{t-1} [\tau (\log(1 - \delta) + \log(\tau) + \log(w_t)) + (1 - \tau) (\log(1 - \tau) + \log(1 - \delta) + \log(w_t))] \\ &= \sum_{t=1}^{\infty} \delta^{t-1} [\log(1 - \delta) + \log(w_t) + \tau \log(\tau) + (1 - \tau) \log(1 - \tau)] = \\ &= \sum_{t=1}^{\infty} \delta^{t-1} [\log(1 - \delta) + \log((R\delta)^{t-1}w_1) + \tau \log(\tau) + (1 - \tau) \log(1 - \tau)] \\ &= \frac{\log(w_1)}{(1 - \delta)} + \frac{1}{(1 - \delta)} \left(\log(1 - \delta) + \tau \log(\tau) + (1 - \tau) \log(1 - \tau) + \frac{\delta \log(R\delta)}{(1 - \delta)} \right). \end{aligned}$$

Turning to $V_2(w_2 | w_1)$, we observe that that c_t^d is already committed in the first period, so that the present value of utility from durable consumption starting in period 2 is

$$(1 - \tau) \sum_{t=2}^{\infty} \delta^{t-2} \log c_t^d = \frac{\delta(1 - \tau)}{(1 - \delta)^2} \log(R\delta) + \frac{(1 - \tau) \log((1 - \tau)w_1)}{1 - \delta}.$$

Since this is a constant, it follows that the agent's problem reduces to maximizing

$$\tau \sum_{t=2}^{\infty} \delta^{t-2} \log x_t$$

subject to $w_2' \equiv w_2 - (1 - \tau)R\delta w_1$ and $w_{t+1}' = R(w_t' - x_t)$. The solution is $x_t = (1 - \delta)w_t'$, so the general form of the maximized present value is

$$\frac{\tau \log(w_t')}{1 - \delta} + K''.$$

Substituting $w_2' \equiv w_2 - (1 - \tau)R\delta w_1$ we see that the maximized present value from non-durable consumption is

$$\frac{\tau \log(w_2 - (1 - \tau)R\delta w_1)}{1 - \delta} + K''.$$

The present value of the utility from durable consumption is

$$\frac{(1 - \tau) \log((1 - \tau)w_1)}{1 - \delta} + K''',$$

where K''' , like the other constants, is independent of w_1 . Thus the overall present value of utility beginning in period 2 is

$$V_2(w_2 | w_1) = \frac{\tau \log(w_2 - (1 - \tau)R\delta w_1)}{1 - \delta} + \frac{(1 - \tau) \log((1 - \tau)R\delta w_1)}{1 - \delta} + K'.$$

Using the fact that $V_2(\delta R w_1 | w_1) = V_1(\delta R w_1)$, we see that

$$\begin{aligned} V_2(R\delta w_1 | w_1) &= \frac{\tau \log(R\delta w_1 - (1 - \tau)R\delta w_1)}{1 - \delta} + \frac{(1 - \tau) \log((1 - \tau)R\delta w_1)}{1 - \delta} + K' \\ &= \frac{\tau \log(\tau) + (1 - \tau) \log(1 - \tau) + \log(R\delta w_1)}{1 - \delta} + K' \\ &= \frac{\log(R\delta w_1)}{(1 - \delta)} + \frac{1}{(1 - \delta)} \left[\log(1 - \delta) + \tau \log(\tau) + (1 - \tau) \log(1 - \tau) + \frac{\delta \log(R\delta)}{(1 - \delta)} \right], \end{aligned}$$

so that

$$K' = \frac{1}{1 - \delta} \left[\log(1 - \delta) + \frac{\delta \log(R\delta)}{(1 - \delta)} \right].$$

We prove Theorem 1 in the general case. The text makes use of the nested CES/log and of the base case.

Theorem 1: (a) For given (x_1, c_1^*) and each $j \in \{S, L\}$ there is a unique solution to

$$\gamma_1^j = \hat{\gamma}_1^j(\gamma_1^j).$$

This solution together with $\tilde{c}_1^j = \min(\hat{c}_1^j(\gamma_1^j)(z_1^j), x_1 + z_1^j)$ and the choice of j that maximizes long-run utility is necessary and sufficient for an optimal solution.

Proof: Consider random unanticipated income \tilde{z}_1^j at the nightclub. If z_1 is the realized income, the short-run self is constrained to consume $c_1 \leq x_1 + z_1$. Period 2 wealth is given by

$$w_2 = R(s_1 + x_1 + z_1 - c_1 - c_1^d) = R(w_1 + z_1 - c_1 - c_1^d).$$

The utility of the long-run self starting in period 2 is given by the solution of the problem without self-control.

Let \tilde{c}_1 be the optimal response to the unanticipated income \tilde{z}_1 . This is a random variable measurable with respect to \tilde{z}_1 . The overall objective of the long-run self is to maximize.

$$\tau \left(Eu(\tilde{c}_1^j, c_1^q) - \bar{g}(x_1, \tilde{c}_1^j, c_1^q) \right) + \frac{\delta}{(1-\delta)} EV(R(w_1 + \tilde{z}_1^j - \tilde{c}_1^j - c_1^d) | w_1) \quad (\text{A.1})$$

Let $\bar{u}_1(x_1, c_1^q) = \max\{Eu(x_1 + \tilde{z}_1^S, c_1^q), Eu(x_1 + \tilde{z}_1^L, c_1^q)\}$ denote the maximum possible utility given c_1^q and the pair of lotteries S, L . Since \bar{u}_1 does not depend on \tilde{c}_1^j , the optimal level of consumption can be determined for each lottery realization by pointwise maximization of (A.1) with respect to $c_1 = c_1^j(z_1^j)$. The value of the marginal cost of self-control is given by

$$\gamma_1 = g'(\bar{u}_1(x_1, c_1^q) - Eu(\tilde{c}_1^j, c_1^q)) = g'(\bar{u}_1(x_1, c_1^q) - \sum_{z_1^j} \Pr(z_1^j) u(z_1^j, c_1^q)). \quad (\text{A.2})$$

First we show that the first order conditions corresponding to optimal consumption for a given choice j have a unique solution. Observe that

$$\frac{d\gamma_1}{dc_1^j(z_1^j)} = -\Pr(z_1^j) u'(c_1^j(z_1^j), c_1^q) g''(\bar{u}_1(x_1, c_1^q) - Eu(\tilde{c}_1^j, c_1^q)) \leq 0.$$

The derivative of (A.1) with respect to $c_1^j = c_1^j(z_1^j)$ evaluated at z_1^j is

$$\tau(1 + \gamma_1) \frac{\partial u(c_1^j, c_1^q)}{\partial c_1^j} - \frac{\delta}{(1-\delta)} \frac{\partial V(R(w_1 + z_1^j - c_1^j - c_1^d) | w_1)}{\partial c_1^j}$$

From this we can compute the second derivative

$$\tau \frac{\partial u(c_1^j, c_1^q)}{\partial c_1^j} \frac{d\gamma_1}{dc_1^j} - \tau \rho (1 + \gamma_1) \frac{\partial^2 u(c_1^j, c_1^q)}{\partial c_1^{j2}} - \frac{\delta}{(1 - \delta)} \frac{\partial^2 V(R(w_1 + z_1^j - c_1^j - c_1^q) | w_1)}{\partial c_1^{j2}} < 0$$

implying that the function is globally concave.

Because the objective function is globally concave with respect to c_1^j , it follows that the unique maximum is given by the solution to the first order condition.

We now show that the conditions in the Theorem are necessary and sufficient for an optimum. Since the problem is one of maximizing a continuous function over a compact space, an optimum exists. To show necessity, Consider any optimum, and suppose the choice is j . Then for any given consumption plan in j the marginal cost of self-control γ_1 is defined by A.2, and the optimal consumption plan must satisfy the first order condition with respect to that γ_1 because our conditions preclude a boundary solution. That is, $\gamma_1^j = \hat{\gamma}_1^j(\gamma_1^j)$ must hold.


Next we show sufficiency. Suppose j, γ_1^j satisfy the conditions of the theorem and that this is not the optimum. That optimum must yield more utility in than choosing $-j$ and any consumption plan in $-j$, so the unique consumption plan that comes from solving $\gamma_1^{-j} = \hat{\gamma}_1^{-j}(\gamma_1^{-j})$. Given that j is chosen, the optimal consumption is the unique solution of the first order condition. On the other hand, if $-j$ was chosen, we could do no better than the consumption plan defined by $\gamma_1^{-j} = \hat{\gamma}_1^{-j}(\gamma_1^{-j})$, and by assumption this is not as good as choosing j .

☑

Computations in the nested CES/logarithmic case: The first order conditions can be written as

$$\begin{aligned}(c_1^j)^\rho &= (c_1^q)^{\rho-1} \frac{\tau(1-\delta)(1+\gamma_1)}{\delta} (w_1 + z_1^j - c_1^j - c_1^d) \\ &= \bar{K} (w_1 + z_1^j - c_1^j - c_1^d)\end{aligned}$$

From the diagram below we can see both the uniqueness of the solution of the first order condition, and also see that the solution is increasing in \bar{K} , that is, decreasing in δ and increasing in γ_1 , and that the solution is increasing in $w_1 + z_1^j$.



We prove Theorem 2 for the base case used in the text.

Theorem 2: Let $\xi \equiv \sup_{c_1} (-c_1 u''(c_1)/u'(c_1)) / \inf_{c_1} (-c_1 u''(c_1)/u'(c_1))$. Then

$$-\frac{V''(w_1)}{V'(w_1)} \leq -\xi \frac{c_1}{w_1} \frac{u''(c_1)}{u'(c_1)}$$

Proof: Let c_t be the optimal plan when initial wealth is w_1 . For $z \geq 0$ define

$$U(z) \equiv \sum_{t=1}^{\infty} \delta^{t-1} u(c_t + (c_t/w_1)z).$$

Notice that the budget constraint implies that

$$\sum_{t=1}^{\infty} R^{-t+1}(c_t / w_1) = 1,$$

so that if z is an increment to wealth the plan $c_t + (c_t / w_1)z$ is feasible and satisfies the budget constraint with equality. This implies first that $V(w_1 + z) \geq U(z)$, and second, from the first order conditions that $V'(w_1) = U'(0)$. Since V is concave, these two facts imply that V is twice differentiable at w_1 and that $-V''(0) \leq -U''(0)$.

To finish the proof, write out

$$\begin{aligned} -U''(0) &= -\sum_{t=1}^{\infty} \delta^{t-1} u''(c_t) (c_t / w_1)^2 \\ &= (1 / w_1) \sum_{t=1}^{\infty} \delta^{t-1} u'(c_t) (c_t / w_1) \{-c_t u''(c_t) / u'(c_t)\} \\ &\leq (\sup_{c_t} \{-c_t u''(c_t) / u'(c_t)\} / w_1) \sum_{t=1}^{\infty} \delta^{t-1} u'(c_t) (c_t / w_1) \\ &= (\sup_{c_t} \{-c_t u''(c_t) / u'(c_t)\} / w_1) U'(0) \end{aligned}$$

Since by the first order conditions of dynamic programming $V'(0) = u'(c_1)$, we may write

$$-V''(0) \leq (\sup_{c_t} \{-c_t u''(c_t) / u'(c_t)\} / w_1) V'(0)$$

and the result follows by noting that

$$\frac{-c_1 u''(c_1) / u'(c_1)}{\inf_{c_t} \{-c_t u''(c_t) / u'(c_t)\}} \geq 1.$$

☑

Theorem 3 considers the case of small anticipated gambles in the nested CES/log case. The text uses only the logarithmic-preferences version of base case.

Theorem 3: *When $\tilde{z}_1^j \ll \hat{z}_1, j = S, L$ the first order conditions necessary for an optimum are*

$$(1 - \delta)x_1 + \delta \frac{E(x_1 + \tilde{z}_1^k)^{1-\rho}}{E(x_1 + \tilde{z}_1^k)^{-\rho}} = (1 - \delta)w_1 \quad (\text{A.3})$$

$$c_1^q = \left(E(x_1 + \tilde{z}_1^k)^{\rho-1} \right)^{1/(\rho-1)} \quad (\text{A.4})$$

where k is the chosen alternative.

Proof: When $\tilde{z}_1^j \ll \hat{z}_1, j = S, L$ all income is spent at the nightclub so $\tilde{c}_1^j = x_1 + \tilde{z}_1^j$, and consequently the gamble most preferred by the short-run self is chosen. There is no self-control cost, so the objective function, given the choice of gamble of gamble j , is

$$Eu(\tilde{c}_1^j, c_1^q) + \frac{\delta}{(1-\delta)} E \log(w_1 - x_1) + K_2. \quad (\text{A.5})$$

The first order condition with respect to c_1^q is

$$\frac{1}{c_1^q} - (c_1^q)^{\rho-2} E(\tilde{c}_1^j)^{1-\rho} = 0$$

which solves to give (A.4).

The first order condition with respect to x_1 is

$$(c_1^q)^{\rho-1} E(x_1 + \tilde{z}_1^k)^{-\rho} - \frac{\delta}{1-\delta} \frac{1}{w_1 - x_1} = 0.$$

Substituting (A.4) and rearranging gives (A.3).

☑