

NBER WORKING PAPER SERIES

DYNAMIC OPTIMIZATION MEETS BUDGETING:
UNRAVELING FINANCIAL COMPLEXITIES

Guidon Fenig
Luba Petersen

Working Paper 32821
<http://www.nber.org/papers/w32821>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2024

This study is financially supported by the Social Sciences and Humanities Research Council of Canada. Insight Development Grant: "Identifying Rules of Thumb in Consumption and Saving Decisions." The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Guidon Fenig and Luba Petersen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Dynamic Optimization Meets Budgeting: Unraveling Financial Complexities
Guidon Fenig and Luba Petersen
NBER Working Paper No. 32821
August 2024
JEL No. C91,E03,E21,E70

ABSTRACT

This paper explores sources of complexity in dynamic optimization, examining how individuals navigate variation in incomes, prices, and returns in ten-period consumption-saving decisions. Our findings reveal that dynamic optimization poses significant challenges, resulting in suboptimal choices even in straightforward scenarios with stable parameters, full information, no uncertainty, and opportunities to learn. These challenges intensify in scenarios involving complexities such as inflation and compounding returns, marked by a pronounced tendency to over-smooth consumption.

Additionally, we introduce a novel budgeting calculator designed to assist with consumption planning and to collect valuable non-choice data on subjects' planning strategies and horizons—an approach not previously utilized in studies of dynamic optimization. We observe significant heterogeneity in planning horizons and ability to optimize given a chosen horizon. Complete planning leads to better performance in more complex scenarios, even when people do not optimally utilize the calculator. However, there is little reoptimization after the first period and participants tend to stick with suboptimal plans for most of their life cycle. The decision to plan is less influenced by the complexity of the economic environment and more by the length of the planning horizon.

Guidon Fenig
Department of Economics
University of Ottawa
gfenig@uottawa.ca

Luba Petersen
Department of Economics
Simon Fraser University
8888 University Drive
Burnaby, BC V5A 1S6
and NBER
lubap@sfu.ca

1 Introduction

In the landscape of financial decision-making, dynamic optimization is a crucial element influencing choices for both individuals and institutions. This domain spans a broad spectrum of essential decisions, including investment strategies, labor markets, environmental policy formation, and especially, the dynamics of saving and consumption. Economic theories often assume that agents are capable of navigating these complex situations effectively. However, empirical evidence suggests a contrasting reality. Many individuals encounter difficulties in planning over time, a challenge emphasized by recent research (Gabaix and Graeber, 2023; Oprea, 2022). But *what* makes dynamic optimization problems complex?

Our paper investigates how people navigate complexities in consumption-saving decisions, pinpointing five primary sources of complexity: shifts in the economic environment, the impact of compounding returns, the length of the planning horizon, computational challenges, and a lack of experience. To explore these elements, we conducted incentivized decision-making tasks where participants undertook a series of consumption-saving exercises over ten-period horizons with full information about all relevant future variables.

In our baseline scenario, ConstantYP, participants faced consumption decisions under conditions of constant and known incomes and prices, and without receiving interest on their savings. In this simplified setting, theoretically optimal behavior dictates spending one's entire income in each period. This design is strategically simple, enabling us to detect patterns of random decision-making and any systematic biases.

We then introduce various facets of complexity within the optimization problem to better understand common pitfalls in decision-making. In the FluctY treatment, we alter participants' income between two predictable levels, high and low. This variation helps us determine if participants can grasp the concept of smoothing their utility from consumption over time as income changes. The FluctP treatment involves changing prices between high and low levels. This treatment variation allows us to observe

whether participants take into account the shape of their induced utility function—that is, how they value each additional unit of consumption—when they decide how much to spend. In the PosIR treatment, we introduce a new element to our baseline setting by introducing positive interest on savings. This enables us to assess how the complexity of decision-making is affected by the presence of compounding returns. Furthermore, across all treatments, participants made spending decisions sequentially over a ten-period horizon. This sequential approach to decision-making allowed us, at the subject-level, to investigate the additional complexity associated with lengthier planning horizons. Specifically, we examine whether participants are more likely to make larger planning errors conditional on their current savings at the start of their decision-making *life-cycle*.

Our experiments reveal that individuals frequently encounter significant challenges in optimizing their lifetime consumption plans, even in scenarios where income and prices are stable and savings yield no returns. These results underscore the inherent difficulty of optimization decisions, even under the simplest environments. In our ConstantYP treatment, the average inexperienced participant managed to attain only 76% of the unconditional optimal solution (i.e the optimal solution that does not account for current savings). This performance marginally improves to 82% by the third repetition of the optimization task, indicating limited learning over time.

Contrary to our expectations, introducing complexity by altering income levels does not hinder participants’ performance. However, the introduction of more complex elements like compounding interest on savings or fluctuating prices significantly reduces participants’ ability to optimize unconditionally. In such complex scenarios, inexperienced participants’ optimization efficiency falls to 66% and 62%, respectively. This decline in performance is persistent, even as participants gain experience through the second and third iterations of the task. Interestingly, shortening the planning horizon is more effective in reducing optimization errors than simply gaining more experience. In the PosIR treatment, for example, participants improve their optimization performance by over 3 percentage points by reducing the planning horizon by just two periods, a level of improvement typically seen after repeating the 10-period task twice.

Furthermore, the nature of the errors varies with the type of complexity introduced. In the PosIR scenario, a common oversight is the undervaluing of early savings, leading to diminished wealth for later consumption. On the other hand, the FluctP scenario reveals a different challenge: while most participants recognize the advantage of spending more when prices are low, their spending often falls short of the optimal level. This is the first study to analyze the complexity of decision-making in the context of inflation in isolation, i.e., when all variables except prices remain constant. It highlights the increased cognitive challenges involved in making optimal decisions under such inflationary conditions.¹

To further investigate the impact of these complexities on decision-making, our study examines the tools individuals commonly use for planning consumption and savings. Typically, individuals rely on basic tools such as simple calculators or their own mathematical skills. This reliance raises an important question: Are the common errors in dynamic optimization primarily due to calculation mistakes? To address this question, our study introduces a second set of treatments where participants have access to a budgeting calculator. This device goes beyond basic calculations, offering tailored assistance for creating comprehensive spending plans. It also tracks all computations, allowing us to analyze how extensively participants plan their spending and whether their planning behavior changes with the economic conditions they face. This innovative approach, focusing on non-choice data, provides a unique lens to assess the depth of people’s dynamic planning and the impact of economic complexity on their short-sightedness. To our knowledge, this study is the first to utilize non-choice calculator data in the context of dynamic optimization. By analyzing this data, we gain insights into the decision-making processes and go beyond just the final decisions made by participants.²

Our analysis into the budgeting calculator’s effectiveness yields mixed results. While

¹While the experiments conducted by [Luhan et al. \(2014\)](#) and [Yamamori et al. \(2018\)](#) have analyzed dynamic optimization under fluctuating prices, they simultaneously alter other macroeconomic variables. In [Luhan et al. \(2014\)](#) experiment, interest rates are positive, and in [Yamamori et al. \(2018\)](#), savings fluctuate in tandem with prices. This complicates the evaluation of how subjects’ behavior changes due to purely inflationary factors.

²An example of how this type of data can be used is found in [Fenig et al. \(2022\)](#), who apply it to study group dynamics in non-linear games.

we anticipated that the calculator would streamline the decision-making process, the actual impact was nuanced. For example, in the ConstantYP treatment, calculator access led to an increase in utility by eight to eleven percentage points compared to the unconditional optimal path, and a 3.4 percent rise in monetary rewards normalized by the maximum potential payoff. In the PosIR treatment, the calculator improved optimization by five to six percentage points, along with a 4.4 percent increase in normalized monetary payoff. However, in FluctY and FluctP, its influence was minimal or indiscernible.³

It should be emphasized that the effectiveness of the budgeting tool in simplifying optimization tasks is contingent upon participants’ proficient utilization. Proficient use of the calculator increases unconditional optimization by 10 to 20 percentage points. It is crucial to highlight that in the PosIR setting—unlike in other treatments where problems are addressable by evaluating one or two periods—calculator use significantly facilitates tackling longer-horizon problems, which are otherwise too challenging to compute manually. Because of these cognitive challenges, one possible heuristic is to compress the planning horizon. Further investigation reveals that a short-span planning model does not fully account for deviations from the optimal consumption path. Nevertheless, even when participants facing positive interest rates formulate full plans—which only about half of them do—the task remains exceedingly challenging. Although there is a notable improvement over participants without calculator access, the results remain suboptimal, as evidenced by their decision-making relative to their planning horizons.

Overall, we find strong support to models positing that complexity hinders dynamic optimization (Gabaix and Graeber, 2023; Woodford, 2019). The most substantial barriers to optimal decision-making include the objective function’s curvature, the impact of compounding returns, and the length of the planning horizon. Our research indicates that while varying income alone does not exacerbate decision-making challenges, the presence of numerical calculations can significantly contribute to the complexity of the task. It is imperative to recognize that the presence of a decision-aid tool like a budget-

³For a more comprehensive analysis, see Appendix E.

ing calculator does not automatically translate to improved outcomes. However, when used effectively, it can be a powerful aid in navigating complex financial decisions.

The structure of the paper is as follows. Section 2 provides a summary of the literature most relevant to this paper. Section 3 presents the theoretical model, describes the different treatments, outlines the experimental design, and lists the main hypotheses. Section 4 offers an overview of the main results. In Section 5, we employ econometric techniques to more formally demonstrate the treatment effects of environmental complexity and the use of the budgeting calculator. In Section 6 we extend our model to allow for short-span planning. Finally, Section 7 concludes the paper.

2 Related Literature

In this subsection, we explore two distinct yet interconnected bodies of research. The first pertains to the complexities inherent in decision-making, while the second focuses on learning-to-optimize experiments within the specific domain of consumption and saving decisions. These two areas of study provide critical insights that inform our understanding of the challenges and behaviors observed in financial decision-making processes.

Many researchers have theorized about the difficulties agents face in dynamically optimizing. A common theme is the inattention to key variables in optimization problems, as detailed in studies by [Schipper \(2014\)](#), [Sims \(2003\)](#), [Maćkowiak and Wiederholt \(2015\)](#), and [Gabaix \(2014\)](#). These works suggest that agents might either neglect important variables or focus on a limited subset due to the costs of processing information. Mis-specified preferences also contribute to what appears to be optimization errors. Contrary to the commonly assumed exponential discounting in many intertemporal planning models, substantial evidence suggests that individuals tend to discount future utility in a hyperbolic manner ([Loewenstein and Prelec, 1992](#); [Laibson, 1997](#); [O'Donoghue and Rabin, 1999](#)).

The literature also highlights the numerical challenges in dynamic optimization. Even the authors of this paper admit to finding the mathematical aspects of extended multi-

period constrained optimization daunting. Simple three-period problems, which might seem elementary, can still pose significant hurdles for many (Gabaix and Graeber, 2023). Emphasizing the importance of this, Lusardi and Wallace (2013) highlight that a firm grasp on quantitative literacy is a cornerstone of financial literacy.

In the context of financial planning over lengthy horizons, Ilut and Valchev (2023) develop a model in which agents, despite perceiving all objective variables relevant to their payoff, encounter subjective uncertainty about optimizing their actions in the given state. This necessitates engaging in costly learning processes to determine the optimal course of action. Our experimental design captures this situation: participants are fully informed of all relevant variables, yet they face challenges in effective problem-solving. Moreover, our novel budgeting calculator enables us to observe their planning process. Our findings align with the dual reasoning model proposed by these authors. Initially, our participants tend to rely on cognitive planning (System 2). However, as they progress through their life cycle, many switch to using intuitive heuristics (System 1). Interestingly, towards the end of their life cycle, as the planning horizon becomes shorter, some participants revert to cognitive planning (System 2).

The complexity of the economic environment and its impact on dynamic optimization is a focal point in the literature. Gabaix and Graeber (2023) investigate a three-period consumption-saving scenario, uncovering that tasks with fluctuating endowments and positive, compounding interest are viewed by subjects as more challenging, leading to increased errors and longer decision times. Consistent with Gabaix and Graeber, we find that complexity in specific environments (PosIR and FluctP) result in poorer optimization. We further show that longer planning horizons exacerbate optimization errors.

Indeed, condensing the planning horizon is a commonly employed strategy to reduce cognitive demands in dynamic optimization tasks. Caliendo and Aadland propose a framework accommodating heterogeneity in planning horizons, where individuals with perfect foresight and those adopting hand-to-mouth strategies represent the extremes. ? analyzes experimental data from an environment similar to our PosIR treatment.⁴ In

⁴Most of the data analyzed is taken from Carbone and Hey (2004).

their study, the horizon spans 25 periods, and the goal is to analyze alternative models that participants may have in mind, allowing for deviations from the optimal consumption path. Thus, they conduct a comparison between models and find that an exponential discounting model in which participants discount future consumption (even if they should not) explains 28% of the data. Furthermore, the rational model—characterized by no discounting, lack of hyperbolic preferences, and a full planning horizon—accounts for 22% of the observations, while the short-span planning horizon explains approximately one fifth of the data. This latter observation is consistent with our findings that short-span planning alone cannot account for the majority of deviations observed in our PosIR treatment.

Dynamic optimization problems are intrinsically complex. The challenges stem from various factors: the extended planning horizons, a wide array of choice variables, the options to borrow or save at differing interest rates, and the unpredictability of future events. Such a diverse range of considerations often leads to decision overload for individuals, nudging them towards simpler heuristics or fallback options in an attempt to manage the overload. [Simon \(1956\)](#) develops a model of 'satisficing' that captures this process by which a decision maker curtails further search. This idea is further developed by [Krosnick \(1991\)](#) who links satisficing to a combination of expected task difficulty, motivation, and ability. [Caplin et al. \(2011\)](#) show evidence that in complex decision tasks, subjects search sequentially and stop searching when they reach their reservation utility. Non-choice data from our Calculator treatments corroborate these observations of incomplete searches and suboptimal decision-making, indicating limited average benefits from tools such as our budgeting calculator. Importantly, we do not observe significant differences in search behavior as *environmental* complexity increases (as we move from ConstantYP to more complex settings). We do, however, find a significant increase in calculator usage as the horizon over which subjects are planning is reduced significantly (in the last two periods of their life-cycle). This further substantiates that the planning horizon is a more salient determinant of complexity for participants than the economic environment variables.

Our analysis of dynamic optimization facilitated by a budgeting calculator contributes to the extensive body of research on household finance puzzles. [D’Acunto and Rossi \(2023\)](#) discuss the success of alternative strategies to mitigate household financial errors: financial advising, nudges, and robo-advising. The design of our budgeting tool falls somewhere in between nudges and robo-advising. The presence of the budgeting calculator nudges participants to at least consider the complete planning horizon when making their consumption decisions. The calculator mitigates the cognitive burdens associated with budgeting calculations but does not advise or encourage participants to make certain plans.⁵ [D’Acunto et al. \(2019\)](#) demonstrate that in the context of portfolio investment, robo-advice helps reduce several behavioral biases, including the disposition effect, trend chasing, and the rank effect. We also find that the introduction of a calculator can significantly improve optimization, with more pronounced effects in simpler environments, such as our ConstantYP treatment. The benefits of our tool in more complex settings involving compounding interest or fluctuating prices are notably less, and suggest a need for additional advising or social learning ([Ballinger et al., 2003](#)).

Overview of Experimental Designs in Dynamic Optimization

The controlled environment of a laboratory is instrumental in unraveling how individuals approach optimization amidst the variability of economic conditions. [Table 1](#) provides a comprehensive summary of the various experimental design differences across the broad *learning-to-optimize* literature ([Duffy, 2016](#)). In most experiments, participants obtain interest from their savings, and there is a prevalent uncertainty about future economic variables, with income being the variable that most frequently changes. The most common planning horizon across these studies is 20 periods, and it is notable that most experimental designs do not involve optimal closed-form solutions. The interplay of complex elements within these experimental settings—such as extended horizons, un-

⁵The one exception is that the calculator informs participants if they have insufficient income over their lifetime to fulfill their consumption plan.

certainty, returns on savings, and income fluctuations—poses a challenge in discerning the impact of different aspects of complexity on optimization behaviors. To address this, our parameterization, as detailed at the bottom of the table, introduces one feature of complexity at a time, allowing for a meticulous examination of each element’s influence on optimization.

Table 1: Related learning-to-optimize experimental literature

Paper	Uncertainty	Interest Rates (IR)	Fluctuating Variable	Borrowing Constraints	No. of Periods	Closed-Form Solution
Hey and Dardanoni (1988)	Yes	>0	Income	Yes	Indefinite	No
Ballinger et al. (2003)	Yes	>0	Income	Yes	60	No
Carbone and Hey (2004)	Yes	>0	Income	Yes	25	No
Brown et al. (2009)	Yes	$=0$	Income	Yes	30	No
Ballinger et al. (2011)	Yes	$=0$	Income	Yes	20	No
Carbone and Duffy (2014)	No	>0	Income	Yes	25	Yes
Luhan et al. (2014)	No	>0	Price and IR	Yes	5	Yes
Carbone and Infante (2015)	Yes	>0	Income	No	15	No
Meissner (2016)	Yes	$=0$	Income	No	20	Yes
Duffy and Li (2019)	No	$=0$	Income	Yes	25	Yes
Yamamori et al. (2018)	No	$=0$	Prices	Yes	20	Yes
Carbone et al. (2019)	Yes	>0	Income	No	15	No
Lu (2022)	No	>0	Income	No	9	No
Miller and Rholes (2023)	Yes	>0	Income	No	20	No
Duffy and Orland (2023)	Yes	$=0$	Income	No	3	No
Gabaix and Graeber (2023)	No	>0	Income	No	3	Yes
Li (2024)	No	>0	Income	Yes	Indefinite	No
This paper						
ConstantYP	No	$=0$	None	Yes	10	Yes
PosIR	No	>0	None	Yes	10	Yes
FluctY	No	$=0$	Income	Yes	10	Yes
FluctP	No	$=0$	Price	Yes	10	Yes

Although direct comparisons among these experiments are difficult due to their differing parameters and information sets, two key patterns consistently emerge in the research. First, there is a noticeable struggle with dynamic planning – individuals often either consume too much or too little, neglecting the long-term impacts of compounding returns. Second, the tendency to optimize decisions appears to enhance as the planning horizon becomes shorter. Echoing this, [Ballinger et al. \(2011\)](#) observes that subjects typically anticipate only up to three periods ahead, underscoring the cognitive limits in longer-term financial planning.

3 Theoretical Model and Experimental Design

The theoretical framework underpinning this study is a standard intertemporal model of life-cycle consumption and savings (See [Modigliani and Brumberg, 1954](#)). Our model is based on a finite-horizon and deterministic framework. Each consumer's goal is to

$$\max_{c_t} \sum_{t=1}^T k \left(\frac{1}{1-\sigma} \right) c_t^{1-\sigma} \quad (1)$$

subject to:

$$p_t c_t + s_t = y_t + (1+r)s_{t-1} \quad (2)$$

$$s_t \geq 0 : \forall t \quad \text{and} \quad s_0 = 0. \quad (3)$$

We assume a concave utility function, specifically a constant relative risk aversion (CRRA) with a parameter σ and a constant k , the variables y_t , s_t , and r represent the consumer's exogenous income, savings, and known and constant interest rate, respectively. The constraint $s_t \geq 0$ implies that borrowing is not allowed.

In this finite horizon model, the consumer faces no uncertainty regarding price and income processes. They make decisions about consumption, denoted as c_t , over T periods, considering their income, y_t , and implicitly decides how much to save at the interest rate r . The consumer pays a price p_t for each unit of consumption.

The optimal consumption path is given by $T - 1$ Euler equations:

$$c_{t+1} = \left(\frac{p_t}{p_{t+1}} (1+r) \right)^{\frac{1}{\sigma}} c_t \quad (4)$$

These Euler equations relate consumption in period t to consumption in period $t + 1$ and must be satisfied within the optimal consumption path. We use the lifetime budget constraint (Equation 5) to derive a system of T equations and T unknowns and then find the optimal consumption level c_t^* for $t \in [1, T]$.

$$\sum_{t=1}^T \frac{p_t c_t}{(1+r)^{t-1}} \leq \sum_{t=1}^T \frac{y_t}{(1+r)^{t-1}} \quad (5)$$

Solving the system of equations also yields the optimal level of consumption for period 1,

$$c_1 = \frac{\sum_{t=1}^T \frac{1}{(1+r)^{t-1}} y_t}{\left(p_1 + p_1^{\frac{1}{\sigma}} \left(\sum_{t=2}^T \left(\frac{p_t}{(1+r)^{t-1}} \right)^{\frac{\sigma-1}{\sigma}} \right) \right)} \quad (6)$$

3.1 Treatments

Our objective in designing the baseline experimental environment was to create a setting that would enable participants to solve the dynamic decision problem without relying on advanced budgeting tools. Therefore, we parameterized our baseline environment, named **ConstantYP**, to maintain income and prices constant throughout the entire consumption horizon. We selected a planning horizon of $T = 10$ periods, which is sufficiently long for participants to plan ahead and also provides ample time for us to observe their learning through stationary repetitions of the environment. Including stationary repetitions is critical in life-cycle experiments because some participants may only understand how to overcome the difficulties of the environment halfway through the sequence, making it impossible for them to adjust their previous consumption choices. Each participant encountered three repetitions of ten periods, and we re-calibrated the environment across repetitions to generate slightly different optimization problems.

We designed and implemented three additional treatments, each introducing a distinct dimension of complexity to understand where and why errors occur in decision-making. The treatments were conducted using a between-subject design. The first treatment, **FluctY**, varies participants' income between two predictable levels, high (y_H) and low (y_L), throughout the horizon. By ensuring that a higher income is received in the first period, we prevent the budget constraint from binding for an optimizing participant, and even for individuals who tend to over-consume compared to the uncon-

ditionally optimal solution. This treatment is particularly insightful, as it evaluates the participants' capability to smooth their consumption over time in the face of fluctuating income. In the second treatment, **PosIR**, the complexity is increased relative to the ConstantYP treatment by introducing a positive interest rate ($r = 0.1$) on savings in every period of each of the three repetitions. This approach allows us to observe the cognitive challenges associated with compounding returns on savings decisions. Lastly, in the **FluctP** treatment, we introduce price fluctuations between two values, high (p_H) and low (p_L), for the entire horizon. This treatment investigates whether participants consider the shape of their utility function—that is, their valuation of each additional unit of consumption—in their spending decisions, as prices vary. These three treatments, by adding different complexity dimensions to the basic optimization setting, help us dissect the nuances of decision-making under varied economic conditions.

To incentivize participants' optimization decisions, we induce a standard constant relative risk aversion (CRRA) per-period utility function:

$$u(c_t) = k \left(\frac{1}{1 - \sigma} \right) c_t^{1 - \sigma}.$$

Here, we set $\sigma = 0.5$ across all treatments. The parameter k is a scaling constant that does not affect the optimal consumption path. In the FluctP treatment, we assign the constant k a value of 2.65, and for the other treatments, we assign it a value of 3.35. This specific adjustment is made to ensure that the optimal life-cycle utility is equalized across all treatments, where the interest rate r is set to zero. The chosen value for σ is intended to create a sufficiently large intertemporal trade-off. Table 2 provides details on the income, price, and interest rate processes for each treatment. It is important to highlight the variations across repetitions: in the second repetition, income is doubled relative to the first. In the third repetition, both income and prices are doubled compared to the first repetition. Although the nominal variables change from the first to the last repetition, the real variables stay consistent, ensuring identical optimal solutions across these repetitions.

Table 2: Parameters of Treatments

			Treatment			
			Constant YP	PosIR	FluctY	FluctP
Repetition 1	y_t	$t \in \{1, 3, 5, 7, 9\}$	1000	1000	1500	1200
		$t \in \{2, 4, 6, 8, 10\}$	1000	1000	500	1200
	p_t	$t \in \{1, 3, 5, 7, 9\}$	100	100	100	150
		$t \in \{2, 4, 6, 8, 10\}$	100	100	100	50
	r		0	0.1	0	0
Repetition 2	y_t	$t \in \{1, 3, 5, 7, 9\}$	2000	2000	3000	2400
		$t \in \{2, 4, 6, 8, 10\}$	2000	2000	3000	2400
	p_t	$t \in \{1, 3, 5, 7, 9\}$	100	100	100	150
		$t \in \{2, 4, 6, 8, 10\}$	100	100	100	50
	r		0	0.1	0	0
Repetition 3	y_t	$t \in \{1, 3, 5, 7, 9\}$	2000	2000	3000	2400
		$t \in \{2, 4, 6, 8, 10\}$	2000	2000	3000	2400
	p_t	$t \in \{1, 3, 5, 7, 9\}$	200	200	200	300
		$t \in \{2, 4, 6, 8, 10\}$	200	200	200	100
	r		0	0.1	0	0

Figure 1 displays the optimal consumption path corresponding to each treatment and repetition. Notably, despite the differences in the income process, the optimal consumption path remains the same for both the Constant YP and FluctY treatments. In this frameworks, consumers maximize utility by consuming 10 units per period in Repetitions 1 and 3, and 20 units per period in Repetition 2. In the PosIR treatment, the optimal consumption path progressively increases, reflecting the absence of discounting. Under the FluctP treatment, the optimal consumption level adjusts in response to price fluctuations, increasing when prices are low and decreasing when they are high, consistent with a bi-periodic pattern. Additionally, Figure 1 illustrates the consumption path $\frac{y_t}{p_t}$, representing an agent's mistaken static optimization approach, where the entire income is consumed each period without considering the intertemporal dynamics.⁶

⁶The exception is the ConstantYP treatment, in which the optimal unconditional consumption path implies not saving.

In the ConstantYP treatment, the unconditional optimal solution corresponds to the ‘Hand-to-Mouth’ (H2M) strategy in the initial period. If a participant opts to save any income during this period, the conditionally optimal consumption path would diverge from the H2M heuristic. Conversely, in all other treatments, a marked divergence is observed between the optimal and H2M consumption paths.⁷




Figure 1: Theoretical Predictions.

We introduce a second between-subject treatment variation to our study by varying participants’ access to a budgeting calculator across each of the four parameterizations. In the **Calc** treatments, participants were provided with the option to utilize a budgeting calculator, enabling them to input hypothetical consumption choices for current and future periods based on the given income, prices, and interest rates. Figure 2(a) shows a screenshot of the budgeting calculator. In this example, participants could input a consumption plan for periods 6-10 and calculate the hypothetical per-period and accumulated points. Participants had unrestricted access to the calculator without any

⁷We analyze this strategy in Appendix D.

mandatory usage requirements. Upon submitting their decisions, participants could review the history of their consumption choices and corresponding savings balances. The software recorded all plans after a participant clicked the calculate button. However, each time subjects submitted their consumption choices, the calculator was cleared.

In contrast, participants in the **NoCalc** treatment lacked access to the budgeting calculator. Figure 2(b) displays a screenshot of the NoCalc interface. Although participants were unable to utilize the tool for financial planning, they retained the ability to review the history of their past choices as well as future incomes, prices, and interest rates. Participants in both treatments could employ the Windows calculator for their decisions; however, these inputs were not recorded.^{8 9}



Figure 2: Screenshots for Treatments with a Budget Calculator and Treatments without a Budget Calculator

⁸It is noteworthy that our NoCalc treatments provide more detailed and readily accessible information compared to standard experiments in consumption smoothing. In these standard experiments, subjects sometimes only see their current income and accumulated savings, based on which they have to make their current consumption decisions.

⁹Household recall of past decisions is shown to be limited [D'Acunto and Weber \(2022\)](#), and households are unable to pay attention and retain all this information over their lifetime [D'Acunto et al. \(2023\)](#). As both of our No-Calculator and Calculator treatments provide participants with detailed information about their past, current, and future individual spending and aggregate variables, we interpret participants' optimization performance as an upper-bound on their ability to solve such problems.

3.2 Experimental Procedures

The study recruited undergraduate and graduate participants from various academic disciplines, resulting in a total of 326 subjects who participated in one of eight treatments. The experiments were conducted in-person at the CRABE Laboratory at Simon Fraser University (SFU) and the SSRL at the University of Saskatchewan (UofS). Our participant sample was nearly evenly split between individuals from SFU and UofS. The experimental sessions were held from July 2017 through May 2018. Table 3 shows the number of participants per treatment. The experiment was computerized and programmed using z-Tree [Fischbacher \(2007\)](#).

Table 3: Number of Subjects per Treatment

Treatment	Budgeting Calculator	Number of Subjects
ConstantYP Calc	Yes	44
ConstantYP No Calc	No	44
PosIR Calc	Yes	48
PosIR No Calc	No	43
FluctY Calc	Yes	48
FluctY No Calc	No	46
FluctP Calc	Yes	45
FluctP No Calc	No	48

We implemented a “between subjects” design by randomly assigning each subject to only one of the eight treatments. At the beginning of each session, the experimenter read aloud the written instructions provided to the participants. Subsequently, participants completed an interactive computerized instruction phase and answered incentivized control questions. These control questions, detailed in Appendix G, ensured that participants understood the main features of the dynamic optimization experiment

and the tools available to them. For instance, to illustrate the principle of diminishing marginal utility, given that the induced utility function is concave, a set of control questions ensured that participants grasp this concept before proceeding with the main tasks of the experiment. In one set of questions, participants were asked to provide the utility levels associated with consumption levels of 1, 2, 9, and 10 units using the Consumption to Utility converter provided in all treatments. They then had to recognize that increasing their consumption from 1 to 2 units results in a larger increase in utility compared to increasing it from 9 to 10 units. Participants earned points based on their performance in the control questions: four points for each correct answer on the first attempt, three points for the second attempt, and two points for the third attempt. No further points were awarded beyond three attempts, and participants were required to answer all questions correctly to proceed.

After completing these questions, participants proceeded to Stage 1. This stage involved three repetitions of one of the dynamic optimization environments shown in Table 2. Figure 3 displays a screenshot of the main computer interface.¹⁰ At the beginning of each period, participants received income (y_t) in tokens and were asked to decide how many units of output they wanted to purchase. To facilitate their decision-making, participants could refer to the plot correlating output with points positioned at the bottom right corner of the screen. Alternatively, they could utilize the Output to Points converter located in the upper right corner of the screen to enhance their understanding of the utility function's properties. In the Calc sessions, participants could also use the budgeting calculator located on the right-hand side of the screen to explore different consumption paths and their implications for own savings and points (utility) before submitting their final choices on the left-hand side of the screen. Our interface design was intended to ensure that all relevant information was readily accessible to participants on a single screen.¹¹

¹⁰There was no time limit for completing any stage of the experiment. In Stage 1, which consisted of three repetitions, subjects spent between 3 and 85 minutes in total.

¹¹Participants who requested it were provided with pen and paper to write down calculations, which could be used at any point during the experiment.




Figure 3: Screenshot of the Experimental Software for Stage 1

In the subsequent three stages of the experiment, we collected data that distinguished participants across three dimensions: risk preferences, financial literacy, and the ability to use backward induction. In Stage 2, risk preferences were assessed through a task adapted from [Eckel and Grossman \(2002\)](#), where participants selected from six distinct 50/50 gambles (refer to Appendix [H.2](#) for a screenshot). This assessment offers a foundational measurement of key attributes influencing various economic decisions, serving as a control variable. In Stage 3, financial literacy was evaluated using a methodology adapted from [Lusardi and Mitchell \(2007\)](#), involving five multiple-choice questions with points awarded for correct first-attempt answers. This evaluation is crucial, as financially literate individuals are better positioned to create optimal budgeting plans, particularly by taking advantage of consumption smoothing and compounding interest rates (refer to Appendix [H.3](#)). In Stage 4, backward induction ability was assessed through the Race to 60 game, proposed by [Bosch-Rosa et al. \(2018\)](#). Participants competed against a computer, selecting numbers from 1 to 10 to reach or exceed 60 first, earning eight points

per win (refer to Appendix H.4 for a screenshot). This assessment is crucial, as dynamic models frequently rely on backward induction for solutions. The experiment concluded with a demographic questionnaire, collecting information on participants' age, gender, education, and employment status.

To determine participants' compensation, we added up the points they earned in all four stages and from correctly answering the control questions. For Stage 1, participants were paid for one randomly selected repetition. Payment was made in Canadian dollars (CAD) at an exchange rate of 25 points = \$1. In addition, participants received a show-up fee of \$7. On average, participants completed the experiment in 40 minutes, and the average payment was \$21.39. Figure 4 provides a summary of the different stages of the experiment.




Figure 4: Sequence of Events in Experimental Sessions.

3.3 Hypotheses

Our analysis centers on two metrics to gauge participants' optimization abilities. The first metric, the Unconditional Optimal Index, quantifies the relative disparity between the actual utility and the utility of the unconditional optimal solution, expressed as a percentage:

$$UncOptIndex_{i,q,r,t} = \gamma^U = 1 - \frac{|U_{i,q,r,t} - U_{q,r,t}^{unc}|}{U_{q,r,t}^{unc}} \quad (7)$$

where $U_{i,q,r,t}$ represents the utility associated with the actual consumption choice made in treatment q , repetition, r , period t by participant i , while $U_{q,r,t}^{unc}$ represents the treatment-repetition-period utility associated with unconditionally optimal consumption. Note that $U_{q,r,t}^{unc}$ is the same for all participants of the same treatment, repetition and period. The index takes a value of 1 when a participant chooses a consumption level that matches the unconditional optimal. The unconditional optimal level of consumption is calculated at the outset of each repetition, prior to any consumption decisions by participants, based on the complete sequence of income, prices, and interest rates for that repetition. The second measure is the Conditional Optimal Index that computes the relative disparity between the actual utility and the utility linked to participant i 's conditional optimal solution, also represented as a percentage:

$$CondOptIndex_{i,q,r,t} = \gamma^C = 1 - \frac{|U_{i,q,r,t} - U_{i,q,r,t}^{cond}|}{U_{i,q,r,t}^{cond}} \quad (8)$$

where $U_{i,q,r,t}^{cond}$ represents the utility associated with subject i 's conditionally optimal consumption in treatment q , repetition r , and period t . The conditional optimal level of consumption is computed for each period as the repetition progresses, contingent on the starting balances of participant i and the forthcoming stream of incomes, prices, and interest rates. The conditional and unconditional optimal solutions coincide in the first period of each repetition.¹²

¹²Note that one of the advantages of our indices is that they account for the loss in payoffs (utility) when deviating from the unconditional and conditional paths, rather than just measuring deviations in

In our experiment, participants have complete information about the lifetime stream of income, prices, and interest rates they will encounter in each sequence. According to standard rational-agent economic theory, agents knowledgeable in solving constrained optimization problems would make spending decisions aligned with the *unconditional* optimal consumption path. Should they temporarily deviate from this optimal path, they would be expected to make *conditionally* optimal decisions thereafter. The features and complexity of the economic environment should not influence ability to optimize. This serves as the basis of our first hypothesis.

Hypothesis 1: (Un)conditional optimal consumption does not differ across environments or a presence of budgeting calculator: $\gamma^{ConstantYP} = \gamma^{FluctY} = \gamma^{FluctP} = \gamma^{PosIR}$.

However, substantial evidence presented in Table 1 demonstrates significant deviations from the (un)conditionally optimal consumption path in spending decisions. Related work by Gabaix and Graeber (2023) and Enke et al. (2023) highlight the role that complexity plays in behavioral biases and optimization errors. They show that people optimize better in environments that are relatively less complex. To identify and rank the complexity of economic environments, we utilize the model from Gabaix and Graeber (2023), which assesses complexity within an intertemporal consumption framework. Their findings indicate that the level of complexity is determined by the number of non-zero terms in a Taylor expansion of the optimal consumption level. Complexity increases with the number of interacting variables when a term is non-zero. Hence, we initiate this analysis by linearizing Equation \ref{eq } through a Taylor expansion:¹³

$$c_1 = \frac{y_1}{p_1} + \frac{1}{Tp_1} \sum_{t=2}^T (y_t - y_1) - \frac{y_1(T-1)}{2\sigma p_1} r + \left(\frac{1-\sigma}{\sigma} \right) \frac{y_1}{Tp_1^2} \sum_{t=2}^T (p_t - p_1) \quad (9)$$

Following Gabaix and Graeber (2023), each element of the problem is assumed to

terms of consumption values.

¹³Appendix A shows the derivation of the Taylor expansion.

have the same level of complexity, $\bar{c}$.¹⁴ We assess that our ConstantYP environment exhibits the least amount of complexity (normalized to zero complexity). The FluctY treatment entails a complexity level of $\bar{c}$, as individuals need to consider the horizon when calculating optimal choices. Finally, FluctP and PosIR are predicted to have a complexity level of $2\bar{c}$, as individuals now need to consider the interaction between the horizon and the elasticity of substitution, σ . This leads to our alternative hypothesis that the relative complexity across environments determines the relative ordering of optimization errors.

Alternative Hypothesis 1: (Un)conditional optimal consumption is ordered based on the complexity of the environment: $\gamma^{ConstantYP} > \gamma^{FluctY} > \gamma^{PosIR} = \gamma^{FluctP}$.

Given that all environments, with the exception of ConstantYP, exhibit inherent complexity, computing the optimal level of consumption becomes a challenging task. It is reasonable to anticipate that choices of individuals with access to a budgeting calculator would be closer to the optimal level of consumption than those without access to it. This motivates our second hypothesis, which posits that the availability of a budgeting tool significantly improves consumption decisions. Access to a budgeting calculator can be particularly helpful for individuals who struggle with computing the optimal level, even in less complex environments.

Hypothesis 2. The budgeting calculator improves (un)conditional optimal consumption: $\gamma^{j,Calc} > \gamma^{j,NoCalc}$ for environment $j \in \{ConstantYP, PosIR, FluctY, FluctP\}$.

By tracking the consumption decisions of participants over multiple repetitions, we can evaluate how experience contributes to enhancements in their dynamic optimization skills. Evidence from [Ballinger et al. \(2003\)](#) suggests that participants learn from

¹⁴For hypothesis formulation and testing purposes, we focus on ordinal ranking rather than cardinal ranking.

experience. Our experimental design can capture this learning clearly. Comparing the first and third repetitions, where changes are made only in nominal terms, the optimal consumption paths remain unaltered.

Hypothesis 3. Deviations of consumption from the (un)conditional optimal consumption paths decrease in later repetitions.

A fundamental distinction in calculating unconditional versus conditional optimal paths lies in the respective planning approaches. For the unconditional path, it is necessary to accurately determine the optimal consumption values for the entire horizon immediately in period 1. This task becomes increasingly challenging with a longer horizon. On the other hand, the conditional optimal path allows subjects to adjust their plans in response to unfolding events as each period progresses, making it easier to stay on or near the optimal path. Consequently, we hypothesize that with a shorter horizon for optimization, subjects are more likely to make consumption decisions that align with the conditional optimal path, as the reduced planning horizon simplifies the task.

Hypothesis 4. Deviations of consumption from the conditional optimal consumption path decrease as the planning horizon becomes shorter.

Our final hypothesis links individual characteristics with dynamic optimization capabilities. Dynamic optimization requires participants not only to perform complex financial calculations but also to contemplate their future decisions. We hypothesize that mathematical proficiency, financial literacy, and backward induction skills will positively correlate with dynamic optimization abilities (see [Lusardi and Wallace, 2013](#)). Participants endowed with these skills are likely better prepared to comprehend and effectively utilize the budgeting tool, thereby achieving greater alignment with the optimal consumption paths.

Hypothesis 5. Dynamic optimization is positively correlated with mathematical training, financial literacy, and backward-induction skills.

4 Overview of experimental results

We begin with an overview of optimization performance and calculator usage across the different treatment.

4.1 Average consumption decisions

Figure 5 presents the mean consumption decisions over time for Repetitions 1 and 3 for each treatment, represented by solid black lines. The gray area indicates the 95% confidence intervals.¹⁵ The optimal consumption path is denoted by red dashed lines. We only include periods 1 to 9 because the consumption decision in Period 10 was trivial; to maximize their current utility, subjects had to spend all the available cash in their bank account.¹⁶

Across different treatments, there are noticeable disparities in participants' optimization abilities. In the ConstantYP treatment, even in the first repetition, the mean participant with access to the budgeting calculator consistently selects a consumption path that is nearly optimal.

In other treatments, deviations from optimality exhibit distinct patterns. In the PosIR treatment, the optimal path involves saving a relatively larger share of income in early periods to take advantage of the compounding return. Participants often underestimate the benefits of early savings in their lifecycle, leading to overconsumption in initial periods and underconsumption in subsequent periods. Their ability to optimize does not appear to improve with experience.¹⁷

¹⁵The individual time series for the consumption choices can be found in Appendix F.

¹⁶To make sure participants had no calculation errors, we displayed a message indicating the exact consumption amount they needed to enter to spend it all. The vast majority followed this guidance. Across all subjects and all repetitions, in 90 percent of the cases, they finished the repetition with no cash remaining in their bank account.

¹⁷We also investigate a scenario where subjects perceive an increase in savings following a simple




Figure 5: Average consumption per period and treatment

In the FluctY treatment, the optimal consumption path remains constant. Optimizing participants should save a portion of their income in odd-numbered periods to offset lower income in even-numbered periods. In Repetition 1, less experienced participants often under-consume in periods with lower income, leading to excessive saving and consumption in Period 9. However, with increased experience, this tendency to oversave diminishes, particularly when participants have access to the budgeting calculator.

In the FluctP treatment, the optimal plan involves setting aside some income during odd-numbered periods when prices are relatively high and utilize their entire wealth interest rate model. This implies that each saved token is believed to grow to an amount represented by $1 + \sum_{s=t}^T r_s$ by the end of a specified time horizon. Nonetheless, the data from PosIR cannot be explained by individuals employing this heuristic.

during even-numbered periods. Although the average participant shows a basic understanding of how consumption should fluctuate with price changes, there is a marked tendency to excessively smooth consumption. This leads to heightened consumption in odd-numbered periods and inadequate consumption in even-numbered periods relative to the unconditionally optimal decision. This behavior remains up to Repetition 3. However, experience leads to notable improvements in optimization, especially for those participants who do not have access to a budgeting calculator.

Finally, regardless of the availability of a budgeting calculator, participants exhibit distinct consumption smoothing behaviors across different treatments. In simpler treatments like ConstantYP and FluctY, we observe a trend towards undersmoothing of consumption among experienced participants. On the other hand, in more complex treatments such as FluctP and PosIR, experienced participants demonstrate a tendency for oversmoothing. While undersmoothing is not necessarily surprising in ConstantYP and FluctY (this is the only direction suboptimal behavior can go in), it was ex-ante unclear how much participants would smooth their consumption in PosIR and FluctP. This pattern suggests that excess consumption smoothing represents a strategic response adopted by experienced individuals when navigating complex economic environments. It resonates with the ‘wait-and-see’ approach, indicating that in scenarios marked by uncertainty or complexity, the preferred strategy is to maintain the current course of action, rather than implementing significant changes in consumption behavior.

4.2 Measuring Optimization Abilities: Unconditional and Conditional Optimal Index

The distributions of γ^U (Unconditional Optimal Index) and γ^C (Conditional Optimal Index) are shown in Figures 6 and 7, respectively. The distributions are broken down by repetition and treatment. We separate the 90 to 95 percentage and 95 to 100 percentage bins to highlight participants’ ability to nearly perfectly optimize, while allowing for the possibility of rounding.




Figure 6: Histogram of UncOptIndex

These histograms reveal three preliminary observations. First, there are significant optimization errors across all of our treatments, even after significant experience. In the simplest environment, where incomes and prices are constant and participants have access to a budgeting calculator, one-quarter of the individuals deviate from the optimal path by more than 10 percent in Repetition 3. And in more complicated environments like PosIR and FluctP, deviations are significantly more substantial and persistent, regardless of calculator accessibility. Surprisingly, this pattern does not change much for deviations from the conditional optimal, as shown in Figure 7.

Second, among those at the upper end of the distribution, participants' abilities to optimize dynamically—both unconditionally and conditionally—enhance with experience, particularly in relatively simple environments such as ConstantYP and FluctY. This improvement is partly attributed to subjects who already had a good understand-




Figure 7: Histogram of CondOptIndex

ing of the problem in the first repetition, as they further improved their optimization skills by the third repetition. Notably, between the first and third repetitions, we observe a significant increase of approximately 15 percentage points in the proportion of participants achieving perfect optimization in ConstantYP, and a 10 percentage point increase in FluctY. However, in more complex environments like PosIR and FluctP, where the optimization challenge is greater, the observed effects are comparatively smaller.

Third, the presence of a budgeting calculator significantly influences optimization in the ConstantYP treatment, more so than in the other treatments, particularly for participants at the top of the distribution. This effect is less obvious for participants in the middle or bottom of the distribution. In the ConstantYP environment, the presence of the budgeting calculator shifts the distribution towards the right, resulting in a 14 percentage point increase in the proportion of participants who achieve perfect optimiza-

tion. However, the benefits provided by the calculator are less pronounced in the other three environments.

4.3 Evaluation of hypotheses

Table 4 presents the mean deviation of the *UncOptIndex* (γ^U) and *CondOptIndex* (γ^C), measured at the session level for each treatment and repetition. Additionally, we report the mean deviations for time spent and the number of calculations attempted in the Calculator treatments. Differences within and across treatments are documented, with corrections applied for multiple hypothesis (List et al., 2019).

We find compelling evidence to reject Hypothesis 1, which predicts that there should be no difference in deviations from both the unconditional and conditional optimal levels of consumption across different environments. Controlling for access to the budget calculator, both γ^U and γ^C are consistently higher in the ConstantYP environment compared to other treatments across most repetitions. Participants also appear better optimize, both unconditionally and conditionally, in the FluctY environment than in either PosIR or FluctP. Notably, PosIR and FluctP exhibit marked similarities, especially when participants do not have access to a budgeting calculator. From a broader perspective, the treatment ordering aligns with Alternative Hypothesis 1 that considers the cognitive complexity present within the economic environment.

Result 1: The ability to (un)conditionally optimize differs across environments:

$$\gamma^{ConstantYP} > \gamma^{FluctY} > \gamma^{PosIR} = \gamma^{FluctP}.$$

Our results provide mixed support for Hypothesis 2, which posits that access to the budgeting calculator significantly improves optimization for both unconditional and conditional consumption decisions. In Repetition 1 and Repetition 3, access to the calculator leads to reductions in mean deviations from the unconditional optimal path by 10.2% and 7.7% in ConstantYP and by 5.8% and 5.9% in PosIR, respectively. Cor-

Table 4: Summary of treatment effects

		Treatment				p-value of two sided t-test					
		Constant YP	Pos IR	Fluct Y	FluctP						
		[1]	[2]	[3]	[4]	[1] vs [2]	[1] vs [3]	[1] vs [4]	[2] vs [3]	[2] vs [4]	[3] vs [4]
Repetition 1											
γ^U	Calculator	0.862	0.718	0.785	0.678	0.000	0.000	0.000	0.000	0.134	0.000
	No Calculator	0.760	0.661	0.778	0.621	0.000	0.718	0.000	0.000	0.286	0.000
	Difference	0.102***	0.058**	0.007	0.057*						
γ^C	Calculator	0.859	0.734	0.839	0.741	0.000	0.038	0.000	0.000	0.758	0.000
	No Calculator	0.780	0.680	0.826	0.693	0.000	0.517	0.000	0.000	0.878	0.000
	Difference	0.079***	0.054**	0.013	0.049***						
Seconds/period	Calculator	35.569	51.951	37.822	60.284	0.264	0.752	0.185	0.332	0.854	0.244
	No Calculator	35.301	38.615	30.600	47.383	0.616	0.647	0.046	0.654	0.666	0.000
	Difference	0.268	13.336	7.221	12.901						
No. Sim/period	Calculator	3.705	5.046	4.060	4.162	0.871	0.966	0.953	0.944	0.958	0.922
Repetition 2											
γ^U	Calculator	0.868	0.701	0.794	0.680	0.000	0.000	0.000	0.000	0.556	0.000
	No Calculator	0.761	0.673	0.822	0.649	0.000	0.000	0.000	0.000	0.536	0.000
	Difference	0.107***	0.028	-0.028	0.031						
γ^C	Calculator	0.869	0.713	0.849	0.736	0.000	0.602	0.000	0.000	0.507	0.000
	No Calculator	0.780	0.677	0.858	0.705	0.000	0.000	0.000	0.000	0.425	0.000
	Difference	0.089***	0.035	-0.008	0.031						
Seconds/period	Calculator	21.792	37.784	25.208	34.988	0.421	0.604	0.038	0.534	0.770	0.215
	No Calculator	17.506	18.736	18.925	31.530	0.856	0.830	0.000	0.942	0.002	0.012
	Difference	4.286	19.048	6.283	3.458						
No. Sim/period	Calculator	1.316	2.621	2.108	1.573	0.510	0.459	0.961	0.986	0.790	0.920
Repetition 3											
γ^U	Calculator	0.897	0.761	0.829	0.703	0.000	0.000	0.000	0.000	0.000	0.000
	No Calculator	0.820	0.702	0.840	0.680	0.000	0.567	0.000	0.000	0.603	0.000
	Difference	0.077***	0.059***	-0.011	0.023						
γ^C	Calculator	0.902	0.766	0.875	0.760	0.000	0.210	0.000	0.000	0.887	0.000
	No Calculator	0.836	0.708	0.888	0.728	0.000	0.000	0.000	0.000	0.471	0.000
	Difference	0.066***	0.059*	-0.013	0.032						
Seconds/period	Calculator	20.146	29.726	22.263	29.771	0.387	0.905	0.269	0.532	0.994	0.432
	No Calculator	13.118	15.518	16.821	21.637	0.527	0.130	0.000	0.482	0.021	0.123
	Difference	7.028	14.209	5.442	8.134						
No. Sim/period	Calculator	1.395	2.075	1.785	1.478	0.910	0.953	0.974	0.976	0.944	0.972
p-value of two sided t-test (comparisons across repetitions)											
		Calculator				No Calculator					
		Constant YP	Pos IR	Fluct Y	FluctP	Constant YP	Pos IR	Fluct Y	FluctP		
γ^u	Rep1 vs Rep 2	0.985	0.741	0.615	0.941	0.998	0.707	0.054	0.693		
	Rep1 vs Rep 3	0.020	0.036	0.112	0.659	0.020	0.191	0.000	0.068		
	Rep2 vs Rep3	0.024	0.000	0.189	0.702	0.024	0.556	0.663	0.568		
γ^c	Rep1 vs Rep 2	0.935	0.650	0.690	0.891	0.986	0.887	0.202	0.842		
	Rep1 vs Rep 3	0.009	0.204	0.131	0.733	0.019	0.615	0.000	0.478		
	Rep2 vs Rep3	0.093	0.011	0.315	0.671	0.005	0.470	0.135	0.737		
Seconds/period	Rep1 vs Rep 2	0.000	0.613	0.144	0.299	0.000	0.113	0.000	0.000		
	Rep1 vs Rep 3	0.000	0.081	0.041	0.138	0.000	0.080	0.000	0.000		
	Rep2 vs Rep3	0.246	0.755	0.769	0.775	0.246	0.696	0.753	0.098		
No. Sim/period	Rep1 vs Rep 2	0.007	0.032	0.012	0.006						
	Rep1 vs Rep 3	0.008	0.010	0.000	0.009						
	Rep2 vs Rep 3	0.965	0.865	0.914	0.823						

Notes: The *p*-values reported are obtained from *t*-tests and have been corrected for multiple hypothesis testing. Significance levels are denoted as follows: **p* < 0.1, ***p* < 0.05, ****p* < 0.01

respondingly, the improvements in the conditional optimal path are 7.9% and 6.6% in ConstantYP, and 5.4% and 5.9% in PosIR for Repetitions 1 and 3. All these differences are statistically significant at the 5% level. However, the evidence indicates that in the FluctY and FluctP treatments, the calculator’s availability does not significantly impact the optimization of either unconditional or conditional consumption.

Result 2: Access to a budgeting tool does not consistently improve dynamic optimization across environments.

Our findings indicate a complex interaction between experience and the use of the budgeting calculator in guiding subjects towards the unconditional and conditional optimal consumption paths. This observation challenges Hypothesis 3, which posits that deviations from these optimal paths should diminish in later repetitions. Our comparative analysis of γ^U and γ^C across different treatments and between repetitions highlights inconsistencies in the effects of the calculator. PosIR Calc is the only treatment that consistently demonstrates improvements in optimization over time. For subjects with access to the calculator in ConstantYP and FluctY, experience does not significantly enhance performance. As previously shown, experience appears to primarily benefit those who were already close to the optimal paths in the first repetition. The limited learning observed over time in calculator treatments may be attributed to subjects gathering most of their learning during the first repetition. While the calculator appears to be useful in exploring various plans and improves learning within a single repetition, this improvement does not seem to carry over to subsequent repetitions.

On the other hand, in situations where participants do not have access to a calculator, there is evidence of learning occurring across multiple repetitions, with the PosIR treatment again being a notable outlier. This finding is significant for experimental designs where subjects accumulate interest on savings. Our results show that without the aid of a calculator, learning from stationary repetitions is limited, emphasizing the crucial role the budgeting tools can play in enhancing learning.

Result 3: Deviations in consumption from the unconditional and conditional optimal consumption paths do not consistently decrease with experience, particularly when subjects have access to a budgeting tool.

Our remaining hypotheses that relate to the planning horizon and demographic characteristics are explored in the following sections where we exploit our experiments’ panel data structure. Before we discuss our final hypothesis tests, we divert readers’ attention to a brief analysis of participants’ usage of the budgeting calculator.

4.4 How Participants Use the Budgeting Calculator

All calculations computed in the budgeting calculator were recorded. This dataset enables an evaluation of how participants engaged with the calculator and the impact of this engagement on decision-making. We observe five distinct patterns in how participants filled in the budgeting calculator in period t .

1. Complete: All entries between t and $T = 10$ are inputted in each calculation.
2. Sequential: Entries and calculations are conducted sequentially. Subjects compute Period t calculations first, followed by Periods t and $t + 1$, and again $t, t + 1, t + 2$, until they reach T .
3. Partial: Less than $T - t + 1$ periods are entered into the calculator.¹⁸
4. Current Period: Only Period t consumption is inputted.
5. No Inputs: No calculations are made.

How would a fully forward-looking, perfectly optimizing agent utilize the calculator? If they understand how to solve the optimization problem, they should not spend time


¹⁸This type of behavior has been modeled by [Caliendo and Aadland \(2007\)](#). In their model, forward-looking individuals solve partial-horizon models. They speculate that such behavior in individuals stems from factors including “...lack of self-control, financial illiteracy, distaste for contemplating old age, or to avoid financial planning costs in an uncertain environment, among others.”

doing any calculations. On the other hand, if they are forward-looking but have difficulty with the computational aspect of the problem, they should submit a complete set of entries before making their spending decisions. We observe participants *complete* the calculator in two ways: either immediately filling in the calculator entirely before pressing the Calculate button, or sequentially with at least $T - t + 1$ calculations, inputting and calculating one period at a time. This latter approach takes more time but provides participants an opportunity to saliently observe the impact of each period’s consumption-saving decision. Participants inputting fewer than $T - t + 1$ periods into the calculator are categorized as *Partial* or *Current-Period*, and are also referred to as *myopic* planners.

It is important to note that the subjects in the Calculator treatments are not required to use the budgeting calculator when making their spending decisions. As shown in Figure 8, the percentage of subjects who do not activate the calculator increases over time. There are several reasons for this trend: (i) subjects do not need to make calculations every period to determine the optimal consumption path, especially in simple environments, (ii) subjects underwent an extensive tutorial and practice period before starting the experiment, (iii) minor changes in the environments between repetitions may allow subjects to rely on past repetitions to approximate the optimal approach, and (iv) subjects may experience fatigue after a few periods. Despite this, roughly one-half of the subjects submit complete sequences of consumption values at the beginning of each repetition. This is consistent with subjects understanding that the environment is deterministic, and in order to calculate the unconditional optimal, they must submit full sequences. Note that if subjects were to attempt to calculate the conditional optimal solution, they may use the calculator in periods 2 to 10, but there is little evidence of that. Furthermore, there is a constant share of subjects using the calculator to enter consumption values for the current period (between 5% and 20% depending on the treatment), with only a few subjects partially or sequentially using the calculator over time.

Considering the notable decline in calculator usage observed after the initial period of each repetition, we assign participants to a specific Calculator Type based solely on

Figure 8: Budgeting Calculator Types Over Time



their calculator usage during the first period. The distribution of types is summarized in Figure 9.

5 Evaluation of optimization

In this section, we formally assess the efficacy of the budgeting calculator in improving consumer optimization. We examine the impacts of the calculator, its usage, learning outcomes, and demographic factors using the following comprehensive random-effects panel regression model for each respective environment:

$$\begin{aligned} \gamma_{i,q,r,t} = & a + b_0 \text{Calculator} + b_1 \text{Calculator} \times \text{CalculatorType}_{i,q,r}^{t=1} \\ & + b_2 V_{i,q,r,t} + b_3 X_i + b_4 \text{Period} + b_5 \text{Repetition} + \mu_i + \text{error}_{i,q,r,t}, \end{aligned}$$




Figure 9: Budgeting calculator usage types, by treatment and repetition

here, $\gamma_{i,q,r,t}$ represents either $UncOptIndex_{i,q,r,t}$ or $CondOptIndex_{i,q,r,t}$. The variable *Calculator* is a dummy variable that takes the value of 1 in treatments where the budgeting calculator is present, and 0 otherwise. Additionally, $CalculatorType_{i,q,r}^{t=1}$ is a categorical variable categorizing participants into one of five calculator usage types, with NoCalc serving as the baseline category.

We include time-varying controls denoted as $V_{i,q,r,t}$, which capture various factors such as the amount of time spent in a given period t (denoted as *Minutes Spent*) and the cumulative count of calculations performed by the participant since the beginning of Repetition 1 (denoted as *No. of Calculations*). It is important to emphasize that participants who chose to use the budgeting calculator had the freedom to either search for the optimal levels or enter any desired consumption values. Therefore, it is crucial to distinguish between participants who effectively used the calculator to find the optimal levels and those who did not. We define a participant as able to find the optimal consumption if, when using the calculator, their *UncOptIndex* and *CondOptIndex* are equal to or greater than 0.90 at any point up to period t within the repetition. Thus, we

introduce the binary variables $\text{Calc } \gamma^U$ and $\text{Calc } \gamma^C$, which take the value of 1 if subjects were able to calculate the optimal consumption levels or 0 otherwise.

Furthermore, time-invariant demographic controls, denoted as X_i , are included to account for participant characteristics unrelated to decisions made in the experiment. These include variables such as math background score (ranging from 1 to 9), financial literacy score (ranging from 0 to 4), backward induction ability (number of wins in Race-to-60, ranging from 0 to 8), gender, age, institution, level of education, experience in previous experiments, risk tolerance, the number of years as a student, and performance on control questions. This last variable measures the accuracy and consistency of responses to a set of incentivized control questions designed to assess understanding of the experimental setup and procedures.

$\text{Repetition} = \{\text{Rep2}, \text{Rep3}\}$ is a vector of dummy variables that captures the effects of learning from experience. $\text{Period} = t$ is a discrete variable that evaluates how the length of the planning horizon affects optimization. Lastly, the subject-specific random effect is denoted by μ_i .

5.1 Impact of Calculator Usage on Dynamic Optimization

The estimated coefficient on the Calculator dummy variable, $\hat{b}_0$, presented in Columns (1), (4), (7), and (10) of Tables 5 and 6 indicates how the presence of the calculator affects optimization. To further evaluate Hypothesis 2, we test whether $\hat{b}_0 > 0$ for each environment. For both unconditional and conditional optimization, we find that $\hat{b}_0$ is very small, ranging between 0-4 p.p. and not statistically different from zero in any environment.

In the initial period, participants can significantly improve their dynamic optimization capabilities by utilizing the calculator. In the simplest environment, ConstantYP, filling in the calculator with only the current period (first period), either partially or completely, improves unconditional optimization by approximately 7 to 11 percentage points (p.p.) and conditional optimization by around 6 to 11 p.p. compared to partici-

pants who do not activate the calculator.

Sequential completion of the calculator resulted in a less sizeable and significant effect. On average, there is only a 4 percentage point improvement for unconditional

Table 5: Dependent Variable: OptUncIndex

	Constant YP			POS IR			Fluct Y			FluctP		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Calculator	0.043 (0.04)			0.006 (0.04)			0.013 (0.04)			0.014 (0.04)		
No. of Calculations		-0.000 (0.00)			-0.001 (0.00)			-0.000 (0.00)			0.001 (0.00)	
Calc γ^U			0.094*** (0.03)			0.194*** (0.03)			0.118*** (0.02)			0.105*** (0.03)
Minutes Spent	-0.018*** (0.00)	-0.018*** (0.00)	-0.019*** (0.00)	-0.009*** (0.00)	-0.006* (0.00)	-0.010*** (0.00)	-0.012*** (0.00)	-0.011** (0.01)	-0.012*** (0.00)	0.002 (0.00)	0.001 (0.00)	0.002 (0.00)
Repetition 2	0.013 (0.01)	0.014 (0.01)	0.014 (0.01)	-0.004 (0.01)	-0.004 (0.01)	-0.005 (0.01)	0.027*** (0.01)	0.027*** (0.01)	0.034*** (0.01)	0.022* (0.01)	0.023* (0.01)	0.019 (0.01)
Repetition 3	0.056*** (0.01)	0.058*** (0.01)	0.059*** (0.01)	0.039*** (0.01)	0.039*** (0.01)	0.033*** (0.01)	0.050*** (0.01)	0.050*** (0.01)	0.055*** (0.01)	0.049*** (0.01)	0.050*** (0.01)	0.044*** (0.01)
Period	-0.004*** (0.00)	-0.004*** (0.00)	-0.004*** (0.00)	-0.003** (0.00)	-0.003** (0.00)	-0.003** (0.00)	-0.006*** (0.00)	-0.006*** (0.00)	-0.006*** (0.00)	0.004** (0.00)	0.004** (0.00)	0.004** (0.00)
Math Level	0.017 (0.01)	0.013 (0.01)	0.013 (0.01)	0.006 (0.01)	0.006 (0.01)	0.003 (0.01)	0.006 (0.01)	0.006 (0.01)	0.008 (0.01)	0.024*** (0.01)	0.024*** (0.01)	0.022** (0.01)
Backward Induction	0.016 (0.01)	0.015 (0.01)	0.013 (0.01)	0.008 (0.01)	0.008 (0.01)	0.003 (0.01)	0.003 (0.02)	0.002 (0.02)	0.004 (0.02)	0.013 (0.01)	0.013 (0.01)	0.012 (0.01)
Financial Literacy	0.015 (0.02)	0.013 (0.02)	0.013 (0.02)	0.003 (0.02)	0.003 (0.02)	0.000 (0.02)	0.008 (0.02)	0.008 (0.02)	0.005 (0.02)	0.000 (0.02)	-0.001 (0.02)	0.001 (0.02)
Gender (1=Female)	-0.026 (0.04)	-0.033 (0.04)	-0.030 (0.04)	-0.032 (0.04)	-0.031 (0.04)	-0.035 (0.03)	0.002 (0.04)	0.001 (0.04)	-0.001 (0.04)	-0.062* (0.04)	-0.063* (0.04)	-0.061* (0.03)
Age	-0.002 (0.00)	-0.002 (0.00)	-0.002 (0.00)	0.004 (0.00)	0.004 (0.00)	0.004 (0.00)	-0.000 (0.00)	-0.000 (0.00)	0.000 (0.00)	-0.005 (0.00)	-0.004 (0.00)	-0.005 (0.00)
Control Questions	0.004 (0.01)	0.005 (0.01)	0.004 (0.01)	0.018*** (0.01)	0.018*** (0.01)	0.016*** (0.01)	0.009 (0.01)	0.009 (0.01)	0.010 (0.01)	-0.006 (0.01)	-0.006 (0.01)	-0.005 (0.01)
<i>Calculator Usage</i>												
Current Period	0.071** (0.03)	0.080*** (0.03)	0.077*** (0.03)	-0.022 (0.03)	-0.020 (0.03)	-0.015 (0.03)	-0.057** (0.03)	-0.056** (0.03)	-0.043 (0.03)	0.003 (0.03)	0.006 (0.03)	0.018 (0.03)
Partial	0.108*** (0.03)	0.115*** (0.03)	0.090*** (0.03)	0.047* (0.03)	0.049* (0.03)	0.061** (0.03)	0.038 (0.03)	0.038 (0.03)	0.045* (0.03)	0.007 (0.03)	0.011 (0.03)	0.013 (0.03)
Sequential	0.040 (0.03)	0.047* (0.03)	-0.009 (0.03)	-0.084*** (0.03)	-0.080** (0.03)	-0.179*** (0.03)	0.011 (0.02)	0.011 (0.02)	-0.054** (0.02)	-0.001 (0.03)	0.002 (0.03)	-0.039 (0.03)
Complete	0.079*** (0.02)	0.087*** (0.02)	0.023 (0.03)	0.090*** (0.03)	0.093*** (0.02)	-0.056* (0.03)	-0.032 (0.02)	-0.033* (0.02)	-0.097*** (0.02)	0.073*** (0.03)	0.076*** (0.02)	0.036 (0.03)
Constant	0.545* (0.31)	0.544* (0.31)	0.589** (0.30)	-0.321 (0.33)	-0.321 (0.33)	-0.178 (0.30)	0.347 (0.34)	0.334 (0.36)	0.286 (0.35)	0.832 (0.51)	0.853* (0.51)	0.813 (0.50)
Observations	2640	2640	2640	2730	2730	2730	2820	2820	2820	2790	2790	2790
Overall R^2	0.12	0.12	0.14	0.08	0.08	0.13	0.05	0.05	0.09	0.08	0.08	0.09

Notes: This table displays the outcomes of a random effects panel regression analysis. The regression models incorporate demographic control variables, including gender, age, prior experimental experience, academic degree, affiliated institution, current year of study, and risk tolerance levels. Additionally, we control for participants' comprehension of the instructions, by adding the participants' performance in the control questions. Standard errors are presented in parentheses. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Dependent Variable: OptCondIndex

	Constant YP			POS IR			FluctY			FluctP		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Calculator	0.037 (0.04)			0.001 (0.04)			-0.000 (0.04)			0.012 (0.03)		
No. of Calculations		-0.001 (0.00)			-0.001 (0.00)			-0.000 (0.00)			0.000 (0.00)	
Calc γ^C			0.095*** (0.02)			0.070*** (0.02)			0.060*** (0.02)			0.091*** (0.02)
Minutes Spent	-0.003 (0.00)	-0.001 (0.00)	-0.003 (0.00)	0.001 (0.00)	0.003 (0.00)	0.000 (0.00)	0.001 (0.00)	0.002 (0.00)	0.001 (0.00)	0.007*** (0.00)	0.006* (0.00)	0.006** (0.00)
Repetition 2	0.015* (0.01)	0.016* (0.01)	0.017** (0.01)	-0.008 (0.01)	-0.008 (0.01)	-0.005 (0.01)	0.024*** (0.01)	0.024*** (0.01)	0.026*** (0.01)	0.011 (0.01)	0.012 (0.01)	0.019 (0.01)
Repetition 3	0.059*** (0.01)	0.060*** (0.01)	0.061*** (0.01)	0.033*** (0.01)	0.033*** (0.01)	0.036*** (0.01)	0.047*** (0.01)	0.047*** (0.01)	0.045*** (0.01)	0.036*** (0.01)	0.037*** (0.01)	0.041*** (0.01)
Period	0.009*** (0.00)	0.009*** (0.00)	0.009*** (0.00)	0.023*** (0.00)	0.023*** (0.00)	0.023*** (0.00)	0.010*** (0.00)	0.010*** (0.00)	0.009*** (0.00)	0.020*** (0.00)	0.020*** (0.00)	0.019*** (0.00)
Math Level	0.017 (0.01)	0.014 (0.01)	0.014 (0.01)	0.005 (0.01)	0.005 (0.01)	0.006 (0.01)	0.004 (0.01)	0.004 (0.01)	0.004 (0.01)	0.019** (0.01)	0.019** (0.01)	0.018** (0.01)
Backward Induction	0.017 (0.01)	0.016 (0.01)	0.012 (0.01)	0.007 (0.01)	0.007 (0.01)	0.006 (0.01)	0.001 (0.01)	0.001 (0.01)	0.002 (0.01)	0.010 (0.01)	0.010 (0.01)	0.010 (0.01)
Financial Literacy	0.018 (0.02)	0.016 (0.02)	0.018 (0.02)	0.002 (0.02)	0.002 (0.02)	0.004 (0.02)	0.010 (0.02)	0.010 (0.02)	0.012 (0.02)	0.003 (0.02)	0.002 (0.02)	0.003 (0.02)
Gender (1=Female)	-0.030 (0.04)	-0.036 (0.04)	-0.036 (0.04)	-0.018 (0.03)	-0.017 (0.03)	-0.016 (0.03)	0.002 (0.03)	0.002 (0.03)	0.003 (0.03)	-0.062** (0.03)	-0.063** (0.03)	-0.058* (0.03)
Age	-0.001 (0.00)	-0.002 (0.00)	-0.002 (0.00)	0.003 (0.00)	0.003 (0.00)	0.003 (0.00)	0.000 (0.00)	0.000 (0.00)	0.001 (0.00)	-0.005 (0.00)	-0.005 (0.00)	-0.005 (0.00)
Control Questions	0.006 (0.01)	0.007 (0.01)	0.006 (0.01)	0.018*** (0.01)	0.018*** (0.01)	0.016** (0.01)	0.010* (0.01)	0.010 (0.01)	0.010* (0.01)	-0.007 (0.01)	-0.007 (0.01)	-0.006 (0.01)
<i>Calculator Usage</i>												
Current Period	0.071*** (0.02)	0.077*** (0.02)	0.061*** (0.02)	0.006 (0.03)	0.006 (0.03)	0.001 (0.03)	-0.043* (0.02)	-0.044* (0.02)	-0.046** (0.02)	-0.000 (0.03)	0.003 (0.03)	0.005 (0.03)
Partial	0.103*** (0.02)	0.110*** (0.02)	0.083*** (0.02)	0.051** (0.03)	0.052** (0.02)	0.044* (0.02)	0.038 (0.02)	0.036 (0.02)	0.031 (0.02)	-0.020 (0.03)	-0.017 (0.03)	0.000 (0.03)
Sequential	0.019 (0.02)	0.025 (0.02)	-0.030 (0.02)	-0.049 (0.03)	-0.047 (0.03)	-0.074** (0.03)	0.015 (0.02)	0.014 (0.02)	-0.002 (0.02)	0.009 (0.03)	0.012 (0.03)	0.017 (0.03)
Complete	0.060*** (0.02)	0.067*** (0.02)	0.008 (0.02)	0.087*** (0.02)	0.089*** (0.02)	0.067*** (0.02)	-0.015 (0.02)	-0.017 (0.02)	-0.046** (0.02)	0.072*** (0.03)	0.075*** (0.02)	0.052** (0.02)
Constant	0.328 (0.28)	0.328 (0.28)	0.404 (0.27)	-0.448 (0.31)	-0.449 (0.31)	-0.403 (0.31)	0.227 (0.29)	0.226 (0.31)	0.217 (0.30)	0.835* (0.45)	0.852* (0.45)	0.802* (0.44)
Observations	2640	2640	2640	2730	2730	2730	2820	2820	2820	2790	2790	2790
Overall R^2	0.16	0.16	0.18	0.15	0.15	0.15	0.07	0.07	0.10	0.12	0.12	0.13

Notes: This table displays the outcomes of a random effects panel regression analysis. The regression models incorporate demographic control variables, including gender, age, prior experimental experience, academic degree, affiliated institution, current year of study, and risk tolerance levels. Additionally, we control for participants' comprehension of the instructions, by adding the participants' performance in the control questions. Standard errors are presented in parentheses. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

optimization and a 2 percentage point improvement for conditional optimization. These improvements are not statistically significant. By contrast, Complete types experience a 8 p.p.(6 p.p.) improvement in (un)conditional optimization.

In the PosIR environment, characterized by compounding interest on savings, notable performance disparities were observed across the different Calculator Types. When participants only complete the consumption decision for the current period, there is no improvement in optimization. However, the Partial and Complete Calculator Types show improvements of 5 and 9 percentage points, respectively, in their optimization. Moreover, imperfectly using the calculator still leads to significant improvements in conditional optimization (Column 6).

Sequential completion does not result in any improvement in conditional optimization, and it can actually have a negative impact on unconditional optimization. On average, participants who sequentially optimize in the first period of PosIR end up 8 percentage points further away from the optimal unconditional solution compared to those who did not use the calculator during the experiment.

The efficacy of the calculator is less pronounced in the FluctY environment. We do not observe consistent improvements in optimization for the Partial, Sequential, and Complete Calculator Types. In comparison to the NoCalc type, participants who only fill in the current period are 6 percentage points further from the optimal unconditional solution and 4 percentage points further from the optimal conditional solution.

Lastly, in the FluctP environment, the calculator benefits are confined to participants who complete it fully. The Complete Calculator Types are 7 percentage points closer to both the optimal unconditional and conditional decisions compared to those who did not use the calculator. Even if a subject fails to find the optimal solution using the calculator, she still improves her conditional optimization by 5 p.p. by filling the calculator in completely.

The specifications presented in columns (3), (6), (9), and (12) of Tables 5 and 6 include dummy variables $\text{Calc } \gamma^U$ and $\text{Calc } \gamma^C$ to control for whether a participant has effectively used the calculator. The estimates demonstrate that effective usage of the calculator significantly improves subjects' performance by 10 to 20 percentage points for the UncOptIndex and 6 and 9 percentage points for the CondOptIndex. These results robustly support a refined version of Hypothesis 2: *effective* utilization of a

Table 7: Dependent Variable: OptCondIndex

	Constant YP		Pos IR		FluctY		FluctP	
	No Calc (1)	Calc (2)	No Calc (3)	Calc (4)	No Calc (5)	Calc (6)	No Calc (7)	Calc (8)
Constant	0.720*** (0.03)	0.815*** (0.02)	0.583*** (0.03)	0.580*** (0.02)	0.774*** (0.02)	0.804*** (0.03)	0.585*** (0.03)	0.657*** (0.02)
Repetition 2	0.000 (0.01)	0.010 (0.01)	-0.003 (0.01)	-0.021* (0.01)	0.032*** (0.01)	0.010 (0.01)	0.012 (0.02)	-0.005 (0.01)
Repetition 3	0.056*** (0.01)	0.043*** (0.01)	0.028* (0.01)	0.032*** (0.01)	0.062*** (0.01)	0.036*** (0.01)	0.035** (0.02)	0.018 (0.01)
Period	0.011*** (0.00)	0.008*** (0.00)	0.018*** (0.00)	0.028*** (0.00)	0.010*** (0.00)	0.006*** (0.00)	0.020*** (0.00)	0.015*** (0.00)
Observations	1320	1320	1290	1440	1380	1440	1440	1350
Overall R^2	0.03	0.02	0.04	0.12	0.03	0.01	0.04	0.03

Notes: This table presents the results of a random effects panel regression. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

calculator significantly reduces deviations from both unconditional and conditional optimal consumption paths. Moreover, the results highlight the importance of carrying out calculations efficiently. Simply entering information into the calculator or doing so incompletely does not guarantee optimization improvements.

5.2 Experience and length of planning horizon

Our data provides strong support for Hypothesis 4, suggesting that conditional optimization improves as participants plan over shorter time horizons. As detailed in Table 7, which disaggregates our results by treatment, we observe significant increases in $\gamma_{i,t}^c$ over time. In the absence of a calculator, participants in the ConstantYP and FluctY treatments exhibit improvements of 1.1 percentage points per period. This increase is even more pronounced in the PosIR and FluctP treatments, at 1.8 p.p. and 2.0 p.p. per period, respectively. These improvements in conditional optimization range from 9 p.p. to 19 p.p. when comparing the first and ninth periods.

Interestingly, the effects are less pronounced for subjects who have access to a calculator, except in the Pos IR treatment, where within-repetition improvement is notable.

In this case, there is an improvement of 25 p.p. from period 1 to period 9. This suggests that the calculator does not significantly aid subjects in correcting mistakes and approaching conditional optimization when the planning horizon is shorter and cognitive load is reduced. A notable exception to this trend is observed in the PosIR treatment, a relevant finding since this type of experiment is commonly conducted. Here, the calculator makes a difference both for within-repetition and between-repetition learning.

Result 4: Deviations from the conditional optimal consumption path diminish as the planning horizon shortens. However, this effect is less pronounced when participants use a calculator. The notable exception to this trend is observed in the Pos IR treatment.

5.3 Demographics

Mathematical training and financial literacy are key factors associated with financial sophistication (Lusardi and Mitchell, 2011; Van Rooij et al., 2012). Under Hypothesis 5, we posited that individuals with mathematical proficiency would excel in sequential calculations and logical reasoning, crucial for formulating optimal consumption plans. Financial literacy complements this by offering a practical framework for applying such calculations to real-world financial decisions, such as assessing future costs and benefits and understanding the time value of money. These combined skills are presumed to significantly aid individuals in solving intertemporal consumption problems, particularly through the cognitive process of backward induction, which we explicitly elicited in our study.

In Tables B.1 and B.2 in Appendix B, we examine each treatment individually to identify whether the presence of the calculator interacts with demographic characteristics. Our results provide little support for Hypothesis 5. Our findings indicate that individual characteristics do not consistently affect subjects' ability to solve intertemporal

consumption problems. No significant effects were observed based on other demographic variables such as gender or age. Higher financial literacy leads to a large and significant improvement in optimization in ConstantYP-NoCalc (8-9 p.p. improvement per correct answer), but not in any other treatment. Stronger backward-induction skills are associated with better *unconditional* optimization in ConstantYP-Calc, but no improvements elsewhere.

Finally, mathematical training has quite a large effect on FluctP-Calc respondents' unconditional and conditional optimization. Unconditional optimization is 3.5 p.p. higher for each additional level of mathematical training. We observe a $3.5 \times 8 = 28$ p.p. difference between those studying in the most mathematically-intensive degrees and those with the least mathematically-intensive degrees (there are 9 categories, so a difference of 8 math levels). In terms of conditional optimization, we observe a $2.8 \times 8 = 22.4$ p.p. difference. Again, this trait does not lead to significantly higher levels of optimization in other treatments.

We are also left intrigued by the finding that mathematical training does not significantly influence optimization in the PosIR treatment. This is at odds with previous literature that underscores the correlation between mathematical skill and the ability to compute compounding returns. An explanation for the inconsistent demographic effects could be the leveling effect of our extensive training phase before the experiment, which included a thorough 30-minute instruction session. This training may have equalized the playing field among subjects with varied backgrounds.

Result 5: Individual and demographic characteristics do not have a consistent effect across economic environments on optimization ability.

6 A Model of Short-Span Planning

A primary reason for sub-optimal consumption plans is short planning horizons. This phenomenon can be rigorously analyzed by refining our existing economic model. In this

revised model, individuals aim to maximize

$$\max_{c_t} \sum_{t=1}^H k \left(\frac{1}{1-\sigma} \right) c_t^{1-\sigma} \quad (10)$$

subject to:

$$p_t c_t + s_t = y_t + (1+r)s_{t-1}, \quad (11)$$

where H is the horizon, it ranges from 1 to 10. In this context, if $H = 10$, then the individual has full foresight (which brings us back to the standard model), and if $H = 1$, the model predicts behavior that aligns with the Hand-to-Mouth heuristic, where individuals consume their income immediately without saving for the future. The implications of different planning horizons are presented in Table C.1 in the Appendix, which outlines the predicted consumption level in Period 1 by treatment and planning horizon.

Our focus here is on PosIR because, in the rest of the treatments, the planning horizon does not significantly affect the optimal level of consumption. In ConstantYP, the optimal consumption level is unaffected by the planning horizon, while in FluctY and FluctP, the primary difference arises between individuals who plan exclusively for the current period and those who consider future periods. Additionally, if H is an even number, the optimal consumption in Period 1 aligns with the solution for a full-span planner.

In PosIR, the model suggests that individuals maximize the benefits of interest rates by extending the planning horizon. Notably, as the planning horizon is longer, consumption in Period 1 diminishes, underscoring the significance of early savings. We leverage our model's predictive power to ascertain a specific planning horizon for each subject based on their observed consumption decisions in the initial period of each repetition. This builds upon the notion that an individual's consumption choice is influenced by their planning horizon and includes a random error component. The error component, ε , follows a normal distribution with mean zero and standard deviation σ . Then, by per-

forming a maximum likelihood estimation, we can calculate the distribution of horizons in the sample and the standard deviation.¹⁹

The estimated parameters are shown in Table C.2 in the appendix. In this table, we note that close to 30% of the observations are not explained by short-span planning model. Among the observations that can be rationalized by the model, approximately 20% of them would correspond to 1-period planners, around 19% would correspond to individuals that plan for 2 or 3 periods in advance, over 10% would consider 8 periods, and close to 10% would use the full horizon.

Additionally, we can cross-validate these model predictions with the actual planning behaviors observed among the subjects in our experiment. This is possible because we have collected detailed planning data for each subject in every period, derived from the usage of the budgeting calculator. To directly compare this data, we examine each subject’s planning horizon in the first period of each repetition by observing how many entries subjects filled out in the calculator during each of the trials. Note that they can fill out as few as one entry and as many as 10 periods (the full horizon).²⁰ Consequently, in Figure 10, we compare the distribution of planning horizons utilized by subjects against the model’s predictions detailed in Table C.2.²¹

¹⁹The details can be found in Appendix C.

²⁰In our analysis, we make two adjustments when examining the data. The first is that we omit intermediate plans for sequential users: if a subject incrementally increases the planning horizon by clicking the *Submit* button each time they enter an additional consumption value, we only retain the final (longest) plan. The second adjustment recognizes that in a single period, subjects may use the calculator to test various potential plans, complicating the analysis as it is unclear which plan they intend to implement. These plans may even differ in their time horizons. To ensure a balanced evaluation, we assign equal weight to each participant using the budgeting calculator in each repetition. This means that every individual’s plan contributes equally to the overall distribution. However, the weight of a specific plan is inversely proportional to the frequency of calculator usage by that subject. This approach mitigates the potential overrepresentation of subjects who are simply exploring the payoff space, preventing their exploratory attempts from disproportionately affecting the distribution.

²¹Note that the first bar represents subjects who did not utilize the calculator for observed planning behavior and in terms of the Partial-Horizon model were not categorized. As a robustness check, we implemented an alternative method by taking consumption values from period 1 in each repetition and comparing them with ten possible consumption values based on the predictions of our model, each corresponding to a different planning horizon. For each horizon, we calculated the squared deviations between the observed consumption figures and the model’s predicted values. The horizon with the smallest squared deviation was identified as the closest to the subject’s observed consumption behavior. If these deviations fell below a certain threshold, the consumption pattern was classified as uncategorized. This approach’s outcomes did not significantly differ from our main analysis findings.




Figure 10: Comparison between the predicted and the observed horizon

Figure 10 demonstrates that the model accurately forecasts the frequency with which a one-period planning horizon is employed. However, it underestimates the use of the full planning horizon; the model predicts its occurrence at around 10%, whereas, in reality, subjects choose this option nearly half of the time. Furthermore, the model’s prediction of a 12% utilization rate for an 8-period planning horizon starkly contrasts with the actual trials, where this horizon is selected in less than 1% of instances. This notable discrepancy prompts a conjectural explanation: subjects who are aware of the benefits of compounding interest rates, but uncertain about the optimal savings amount, might resort to a straightforward heuristic. Such a heuristic might involve saving half of their income, yielding a consumption level of 5 units, closely aligning with the model’s prediction of 5.13 units for an 8-period planner, as detailed in Table C.1.

7 Discussion

Our study represents one of the first attempts to understand the role of environmental complexity in dynamic optimization, focusing on consumption smoothing. We uncover that different environments pose varying levels of challenges, some more demanding than others. Particularly, while the complexities associated with compounding interest rates on savings or fluctuating income are well-acknowledged, the challenges related

to fluctuating prices have been less examined. Our research not only fills this gap but also presents new empirical evidence demonstrating that many individuals struggle with planning over extended periods, a fundamental yet often overlooked assumption of economic theory.

In exploring these complexities, our experiments reveal that individuals frequently encounter significant hurdles in achieving optimal outcomes, even in seemingly straightforward scenarios. These challenges arise not only from the variability of economic conditions but also from cognitive limitations and numerical difficulties inherent in dynamic decision-making. Our findings resonate with existing literature on complexity, such as the works of [Oprea \(2022\)](#), [Enke and Graeber \(2019\)](#), and [Gabaix and Graeber \(2023\)](#), which conceptualize the mind as a cognitive economy faced with numerous decisions.

Another key aspect of our study is the potential role of a meticulously designed budgeting calculator in aiding decision-making. Our results indicate that the mere availability of such a tool does not guarantee its effective use. The effectiveness of a budgeting tool largely depends on how individuals choose to utilize it. This observation is crucial, as it highlights that providing tools for financial planning is beneficial, but their impact hinges on the user’s engagement and understanding. Furthermore, our analysis of non-choice data from the budget calculator provides deeper insights into individuals’ strategies, preferences, and decision-making horizons. The integration of robo-advisory features into this budgeting calculator could enhance its utility. Robo-advisors, by providing automated, data-driven guidance and personalized financial recommendations, can serve as an active form of nudging, potentially increasing the effectiveness of such tools. Such integration could transform the calculator from a passive instrument into a more interactive agent in financial planning, nudging users towards optimal financial behaviors tailored to their unique data and preferences.

Crucially, our paper makes a significant contribution by directly examining individuals’ capacity for planning, a core assumption in economic theory often taken for granted. Standard economic models assume that agents can plan effectively, even for

long horizons, which is essential for making optimal present choices. However, empirical evidence supporting this assumption, particularly from laboratory experiments, is scarce.²² Our work stands as one of the first to provide concrete evidence of individuals’ planning abilities—or the lack thereof—by directly observing their plans. This approach differs markedly from indirect methods using decision trees (see [Hey 2002](#)), which infer planning from decision outcomes. Expanding upon this concept, we can even compare observed planning-horizons to implied horizons from a short-span horizon model. This analysis leads us to rule out the short-span model in positive interest rate environments, pointing towards a broader issue in achieving optimal paths than just the incapacity for long-term planning.

It is important to emphasize that real-world settings present even greater challenges for dynamic optimization. The unpredictability of real life, combined with unexpected events and incomplete information, complicates decision-making. Individuals often pursue multiple goals and are influenced by behavioral biases like overconfidence, loss aversion, and procrastination. Therefore, the deviations from the optimal path observed in our controlled experiments should be viewed as a lower bound. In more complex real-world scenarios, these deviations are expected to increase, highlighting the necessity for strong financial planning tools and improved financial literacy to help individuals navigate these challenges effectively.

Looking forward, our study opens several avenues for future research. Investigating dynamic optimization under diverse conditions such as uncertainty, social learning influences, expected liquidity constraints, and current binding constraints could enhance our understanding of decision-making complexities. Further studies on the effects of introducing and removing budget calculators and the sustained impact of financial literacy initiatives will provide deeper insights into how these tools and educational interventions affect consumer behavior.

In terms of policy implications, our findings on suboptimal decision-making due to

²²One exception is planning in the context of present bias and procrastination (see [Della Vigna and Malmendier 2006](#)).

cognitive complexity, especially under fluctuating prices, suggest an added layer of welfare loss beyond traditional concerns like menu costs and relative price dispersion in inflationary contexts. These insights highlight the potential for cognitive errors to exacerbate welfare losses in realistic settings, thus underscoring the importance of simplifying financial decision environments and improving financial literacy.

References

- BALLINGER, T., M. PALUMBO, AND N. WILCOX (2003): “Precautionary Savings and Social Learning Across Generations: An Experiment,” *The Economic Journal*, 113, 920–947.
- BALLINGER, T. P., E. HUDSON, L. KARKOVIATA, AND N. T. WILCOX (2011): “Saving Behavior and Cognitive Abilities,” *Experimental Economics*, 14, 349–374.
- BOSCH-ROSA, C., T. MEISSNER, AND A. BOSCH-DOMÈNECH (2018): “Cognitive bubbles,” *Experimental Economics*, 21, 132–153.
- BROWN, A., E. CHUA, AND C. CAMERER (2009): “Learning and Visceral Temptation in Dynamic Saving Experiments,” *Quarterly Journal of Economics*, 124, 197–231.
- CALIENDO, F. AND D. AADLAND (2007): “Short-term planning and the life-cycle consumption puzzle,” *Journal of Economic Dynamics and Control*, 31, 1392–1415.
- CAPLIN, A., M. DEAN, AND D. MARTIN (2011): “Search and satisficing,” *American Economic Review*, 101, 2899–2922.
- CARBONE, E. AND J. DUFFY (2014): “Lifecycle Consumption Plans, Social Learning and External Habits: Experimental Evidence,” *Journal of Economic Behavior & Organization*, 106, 413–427.
- CARBONE, E., K. GEORGALOS, AND G. INFANTE (2019): “Individual vs. Group Decision-Making: An Experiment on Dynamic Choice under Risk and Ambiguity,” *Theory and decision*, 87, 87–122.
- CARBONE, E. AND J. HEY (2004): “The Effect of Unemployment on Consumption: An Experimental Analysis,” *The Economic Journal*, 114, 660–683.
- CARBONE, E. AND G. INFANTE (2015): “Are Groups Better Planners than Individuals? An Experimental Analysis,” *Journal of Behavioral and Experimental Economics*, 57, 112–119.

- D’ACUNTO, F. AND A. G. ROSSI (2023): “Robo-advice: Transforming households into rational economic agents,” *Annual Review of Financial Economics*, 15, 543–563.
- D’ACUNTO, F., A. G. ROSSI, AND M. WEBER (2023): “Crowdsourcing peer information to change spending behavior,” *Chicago Booth Research Paper*.
- DELLA VIGNA, S. AND U. MALMENDIER (2006): “Paying not to go to the gym,” *american economic Review*, 96, 694–719.
- DUFFY, J. (2016): “Macroeconomics: A Survey of Laboratory Research,” *Handbook of Experimental Economics, Volume 2*.
- DUFFY, J. AND Y. LI (2019): “Lifecycle consumption under different income profiles: Evidence and theory,” *Journal of Economic Dynamics and Control*, 104, 74–94.
- DUFFY, J. AND A. ORLAND (2023): “Liquidity Constraints, Income Variance, and Buffer Stock Savings: Experimental Evidence,” Working Paper.
- D’ACUNTO, F., N. PRABHALA, AND A. G. ROSSI (2019): “The promises and pitfalls of robo-advising,” *The Review of Financial Studies*, 32, 1983–2020.
- D’ACUNTO, F. AND M. WEBER (2022): “Memory and beliefs: Evidence from the field,” *Georgetown University and University of Chicago Working Paper*.
- ECKEL, C. C. AND P. J. GROSSMAN (2002): “Sex Differences and Statistical Stereotyping in Attitudes Toward Financial Risk,” *Evolution and Human Behavior*, 23, 281–295.
- ENKE, B. AND T. GRAEBER (2019): “Cognitive uncertainty,” Tech. rep., National Bureau of Economic Research.
- ENKE, B., T. GRAEBER, AND R. OPREA (2023): “Complexity and Time,” Tech. rep., National Bureau of Economic Research.
- FENIG, G., G. GALLIPOLI, AND Y. HALEVY (2022): “Piercing the ‘Payoff Function’ Veil: Tracing Beliefs and Motives,” Working Paper.

- FISCHBACHER, U. (2007): “z-Tree: Zurich Toolbox for Ready-Made Economic Experiments,” *Experimental Economics*, 10, 171–78.
- GABAIX, X. (2014): “A sparsity-based model of bounded rationality,” *The Quarterly Journal of Economics*, 129, 1661–1710.
- GABAIX, X. AND T. GRAEBER (2023): “The Complexity of Economic Decisions,” *Available at SSRN*.
- HEY, J. AND V. DARDANONI (1988): “Optimal Consumption Under Uncertainty: An Experimental Investigation,” *The Economic Journal*, 98, 105–116.
- HEY, J. D. (2002): “Experimental economics and the theory of decision making under risk and uncertainty,” *The Geneva Papers on Risk and Insurance Theory*, 27, 5–21.
- ILUT, C. AND R. VALCHEV (2023): “Economic agents as imperfect problem solvers,” *The Quarterly Journal of Economics*, 138, 313–362.
- KROSNICK, J. A. (1991): “Response strategies for coping with the cognitive demands of attitude measures in surveys,” *Applied cognitive psychology*, 5, 213–236.
- LAIBSON, D. (1997): “Golden eggs and hyperbolic discounting,” *The Quarterly Journal of Economics*, 112, 443–478.
- LI, Y. (2024): “Do Tax Deferred Accounts Improve Lifecycle Savings? Experimental Evidence,” *Review of Economics and Statistics*, forthcoming.
- LIST, J. A., A. M. SHAIKH, AND Y. XU (2019): “Multiple hypothesis testing in experimental economics,” *Experimental Economics*, 22, 773–793.
- LOEWENSTEIN, G. AND D. PRELEC (1992): “Anomalies in intertemporal choice: Evidence and an interpretation,” *The Quarterly Journal of Economics*, 107, 573–597.
- LU, K. (2022): “Overreaction to capital taxation in saving decisions,” *Journal of Economic Dynamics and Control*, 144, 104541.

- LUHAN, W. J., M. W. ROOS, AND J. SCHARLER (2014): “An Experiment on Consumption Responses to Future Prices and Interest Rates,” in *Experiments in macroeconomics*, Emerald Group Publishing Limited, 139–166.
- LUSARDI, A. AND O. MITCHELL (2007): “Financial Literacy and Retirement Planning: New Evidence from the Rand American Life Panel,” NBER Working Paper.
- LUSARDI, A. AND O. S. MITCHELL (2011): “Financial literacy and planning: Implications for retirement wellbeing,” .
- LUSARDI, A. AND D. WALLACE (2013): “Financial literacy and quantitative reasoning in the high school and college classroom,” *Numeracy*, 6, 1.
- MAĆKOWIAK, B. AND M. WIEDERHOLT (2015): “Business cycle dynamics under rational inattention,” *The Review of Economic Studies*, 82, 1502–1532.
- MEISSNER, T. (2016): “Intertemporal Consumption and Debt Aversion: An Experimental Study,” *Experimental Economics*, 19, 281–298.
- MILLER, L. AND R. RHOLES (2023): “Joint vs. Individual performance in a dynamic choice problem,” *Journal of Economic Behavior & Organization*, 212, 897–934.
- MODIGLIANI, F. AND R. BRUMBERG (1954): “Utility Analysis and the Consumption Function: An Interpretation of Cross Section Data,” in *Post Keynesian Economics*, ed. by K.K.Kurihara, Ruthers University Press New Brunswick New Jersey, vol. 1, 388–436.
- O’DONOGHUE, T. AND M. RABIN (1999): “Doing it now or later,” *American Economic Review*, 89, 103–124.
- OPREA, R. (2022): “Simplicity equivalents,” *Working Paper*.
- SCHIPPER, B. C. (2014): “Unawareness—a gentle introduction to both the literature and the special issue,” *Mathematical Social Sciences*, 70, 1–9.

- SIMON, H. A. (1956): “Rational choice and the structure of the environment.” *Psychological review*, 63.
- SIMS, C. A. (2003): “Implications of rational inattention,” *Journal of monetary Economics*, 50, 665–690.
- VAN ROOIJ, M. C., A. LUSARDI, AND R. J. ALESSIE (2012): “Financial literacy, retirement planning and household wealth,” *The Economic Journal*, 122, 449–478.
- WOODFORD, M. (2019): “Monetary policy analysis when planning horizons are finite,” *NBER macroeconomics annual*, 33, 1–50.
- YAMAMORI, T., K. IWATA, AND A. OGAWA (2018): “Does money illusion matter in intertemporal decision making?” *Journal of Economic Behavior & Organization*, 145, 465–473.

A Taylor Expansion for Optimal Choice

In this section, we apply a Taylor expansion to the optimal consumption choice for period 1, as specified in Equation 6. This analysis adheres to the methodology delineated by Gabaix and Graeber (2023) in their study of complexity. Let $\hat{y}_t = y_t - y_1$ and $\hat{p}_t = p_t - p_1$, and assume that $\hat{y}_t$ and $\hat{p}_t$ are negligibly small. Consequently, we can reformulate Equation 6.

$$c_1 = \frac{\sum_{t=1}^T \frac{y_1}{(1+r)^{t-1}} + \sum_{t=2}^T \frac{\hat{y}_t}{(1+r)^{t-1}}}{\left(p_1 + p_1^{\frac{1}{\sigma}} \left(\sum_{t=2}^T \left(\frac{\hat{p}_t + p_1}{(1+r)^{t-1}} \right)^{\frac{\sigma-1}{\sigma}} \right) \right)} \quad (12)$$

We begin by evaluating c_1 when $\hat{y}_t = 0$, $r = 0$ and $\hat{p}_t = 0$

$$c_1 = \frac{T y_1}{p_1 + p_1^{\frac{1}{\sigma}} \left((T-1) p_1^{\frac{\sigma-1}{\sigma}} \right)} = \frac{y_1}{p_1} \quad (13)$$

Then we calculate $\frac{\partial c_1}{\partial \hat{y}_t}$ and evaluate it when $\hat{y}_t = 0$, $r = 0$ and $\hat{p}_t = 0$

$$\begin{aligned} \frac{\partial c_1}{\partial \hat{y}_t} &= \frac{\sum_{t=1}^T \left[\frac{1}{(1+r)^{t-1}} \right]}{\left(p_1 + p_1^{\frac{1}{\sigma}} \left(\sum_{t=2}^T \left(\frac{\hat{p}_t + p_1}{(1+r)^{t-1}} \right)^{\frac{\sigma-1}{\sigma}} \right) \right)} \\ &= \frac{1}{p_1 + p_1^{\frac{1}{\sigma}} \left(p_1^{\frac{\sigma-1}{\sigma}} (T-1) \right)} \\ &= \frac{1}{T p_1} \end{aligned}$$

We then calculate $\frac{\partial c_1}{\partial r}$ and evaluate it when $\hat{y}_t = 0$, $r = 0$ and $\hat{p}_t = 0$

$$\begin{aligned}
\frac{\partial c_1}{\partial r} &= - \frac{\left(y_1 \sum_{t=1}^{T-1} t \right) T p_1 + \left(\left(\frac{1-\sigma}{\sigma} \right) p_1 \sum_{t=1}^{T-1} t \right) T y_1}{T^2 p_1^2} = \\
&= - \frac{y_1 \sum_{t=1}^{T-1} t}{T p_1 \sigma} \\
&= \frac{\frac{-y_1(T-1)(T)}{2}}{T p_1 \sigma} \\
&= - \frac{y_1 (T-1)}{2 \sigma p_1}
\end{aligned}$$

Finally we calculate $\frac{\partial c_1}{\partial \hat{p}_t}$ and evaluate it when $\hat{y}_t = 0$, $r = 0$ and $\hat{p}_t = 0$

$$\begin{aligned}
\frac{\partial c_1}{\partial \hat{p}_t} &= - \frac{\left(\sum_{t=1}^T \frac{1}{(1+r)^{t-1}} (\hat{y}_t + y_1) \right) p_1^{\frac{1}{\sigma}} \left(\left(\frac{\sigma-1}{\sigma} \right) \left(\frac{\hat{p}_t + p_1}{(1+r)} \right)^{-\frac{1}{\sigma}} \right)}{\left(p_1 + p_1^{\frac{1}{\sigma}} \left(\sum_{t=2}^T \left(\frac{\hat{p}_t + p_1}{(1+r)^{t-1}} \right)^{\frac{\sigma-1}{\sigma}} \right) \right)^2} \\
&= - \frac{T y_1 \left(\frac{\sigma-1}{\sigma} \right)}{T^2 p_1^2} \\
&= \left(\frac{1-\sigma}{\sigma} \right) \frac{y_1}{T p_1^2}
\end{aligned}$$

Thus, the Taylor expansion is

$$c_1 = \frac{y_1}{p_1} + \frac{1}{T p_1} \sum_{t=2}^T \hat{y}_t - \frac{y_1 (T-1)}{2 \sigma p_1} r + \left(\frac{1-\sigma}{\sigma} \right) \frac{y_1}{T p_1^2} \sum_{t=2}^T \hat{p}_t \quad (14)$$

B Additional Regressions

Table B.1: Dependent Variable: OptUncIndex

	Constant YP			Pos IR			FluctY			FluctP		
	No Calc (1)	Calc (2)	Calc (3)	No Calc (4)	Calc (5)	Calc (6)	No Calc (7)	Calc (8)	Calc (9)	No Calc (10)	Calc (11)	Calc (12)
No of Calculations			-0.001 (0.00)			-0.001 (0.00)			-0.001 (0.00)			0.001 (0.00)
Calc γ^U			0.093*** (0.02)			0.213*** (0.02)			0.115*** (0.02)			0.105*** (0.02)
Minutes Spent	-0.037*** (0.01)	-0.012*** (0.00)	-0.010* (0.00)	-0.008 (0.01)	-0.009*** (0.00)	-0.005 (0.00)	-0.011 (0.01)	-0.014*** (0.00)	-0.007 (0.01)	0.011 (0.01)	-0.000 (0.00)	-0.002 (0.00)
Repetition 2	-0.010 (0.02)	0.036*** (0.01)	0.038*** (0.01)	0.010 (0.01)	-0.019 (0.01)	-0.019 (0.01)	0.042*** (0.01)	0.007 (0.01)	0.023* (0.01)	0.031* (0.02)	0.005 (0.02)	-0.001 (0.02)
Repetition 3	0.047*** (0.02)	0.065*** (0.01)	0.069*** (0.01)	0.038*** (0.01)	0.037*** (0.01)	0.025* (0.01)	0.059*** (0.01)	0.035*** (0.01)	0.047*** (0.01)	0.063*** (0.02)	0.029 (0.02)	0.018 (0.02)
Period	-0.010*** (0.00)	0.001 (0.00)	0.001 (0.00)	-0.006*** (0.00)	0.000 (0.00)	-0.000 (0.00)	-0.003* (0.00)	-0.009*** (0.00)	-0.009*** (0.00)	0.006** (0.00)	0.002 (0.00)	0.002 (0.00)
Math Level	0.040 (0.02)	0.009 (0.01)	0.007 (0.01)	0.007 (0.02)	-0.015 (0.01)	-0.016* (0.01)	-0.005 (0.02)	0.004 (0.01)	0.005 (0.01)	0.015 (0.01)	0.041*** (0.01)	0.035** (0.01)
Backward Induction	-0.001 (0.02)	0.042** (0.02)	0.036* (0.02)	-0.015 (0.03)	0.015 (0.01)	0.005 (0.01)	0.024 (0.02)	-0.004 (0.03)	-0.007 (0.02)	-0.002 (0.03)	0.027 (0.02)	0.028 (0.02)
Financial Literacy	0.092** (0.04)	-0.018 (0.02)	-0.017 (0.02)	0.022 (0.03)	-0.013 (0.03)	-0.020 (0.02)	0.034 (0.03)	-0.012 (0.04)	-0.013 (0.04)	0.021 (0.03)	-0.012 (0.02)	-0.011 (0.02)
Control Questions	-0.003 (0.01)	0.005 (0.01)	0.003 (0.01)	0.013 (0.01)	0.021** (0.01)	0.020*** (0.01)	0.007 (0.01)	0.012 (0.01)	0.014 (0.01)	0.006 (0.02)	0.002 (0.01)	-0.001 (0.01)
<i>Calculator Usage</i>												
Current Period		0.074*** (0.02)	0.075*** (0.02)		-0.028 (0.03)	-0.006 (0.03)		-0.068** (0.03)	-0.049* (0.03)		-0.002 (0.03)	0.013 (0.03)
Partial		0.122*** (0.02)	0.102*** (0.03)		0.035 (0.03)	0.052** (0.03)		0.016 (0.03)	0.038 (0.03)		-0.007 (0.03)	-0.003 (0.03)
Sequential		0.051** (0.02)	-0.000 (0.03)		-0.089*** (0.03)	-0.176*** (0.03)		-0.005 (0.02)	-0.056** (0.02)		-0.022 (0.03)	-0.063* (0.03)
Complete		0.088*** (0.02)	0.029 (0.02)		0.081*** (0.02)	-0.086*** (0.03)		-0.048** (0.02)	-0.099*** (0.03)		0.053* (0.03)	0.013 (0.03)
Constant	0.677 (0.48)	0.601 (0.44)	0.662* (0.40)	-0.122 (0.73)	-0.360 (0.39)	-0.203 (0.24)	0.267 (0.51)	0.432 (0.53)	0.266 (0.52)	0.311 (1.10)	0.358 (0.68)	0.526 (0.67)
Observations	1320	1320	1320	1290	1440	1440	1380	1440	1440	1440	1350	1350
Overall R^2	0.15	0.16	0.21	0.05	0.14	0.26	0.10	0.05	0.14	0.07	0.14	0.16

Notes: This table displays the outcomes of a random effects panel regression analysis. The regression models incorporate demographic control variables, including gender, age, prior experimental experience, academic degree, affiliated institution, current year of study, and risk tolerance levels. Additionally, we control for participants' comprehension of the instructions, by adding the participants' performance in the control questions. Standard errors are presented in parentheses. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.2: Dependent Variable: OptCondIndex

	Constant YP			Pos IR			FluctY			FluctP		
	No Calc (1)	Calc (2)	Calc (3)	No Calc (4)	Calc (5)	Calc (6)	No Calc (7)	Calc (8)	Calc (9)	No Calc (10)	Calc (11)	Calc (12)
No of Calculations			-0.001 (0.00)			0.001 (0.00)			-0.001 (0.00)			0.001 (0.00)
Calc γ^C			0.112*** (0.02)			0.052*** (0.01)			0.032* (0.02)			0.057*** (0.02)
Minutes Spent	0.002 (0.01)	-0.005* (0.00)	-0.004 (0.00)	0.008 (0.01)	0.000 (0.00)	-0.003 (0.00)	-0.006 (0.01)	-0.005 (0.00)	-0.001 (0.01)	0.019*** (0.01)	0.002 (0.00)	-0.001 (0.00)
Repetition 2	0.001 (0.01)	0.034*** (0.01)	0.062*** (0.01)	-0.000 (0.01)	-0.014 (0.01)	-0.005 (0.01)	0.031*** (0.01)	0.008 (0.01)	0.014 (0.01)	0.017 (0.02)	0.005 (0.02)	0.017 (0.02)
Repetition 3	0.057*** (0.01)	0.067*** (0.01)	0.080*** (0.01)	0.031** (0.01)	0.037*** (0.01)	0.044*** (0.01)	0.061*** (0.01)	0.030** (0.01)	0.030** (0.01)	0.043*** (0.02)	0.030* (0.02)	0.041** (0.02)
Period	0.011*** (0.00)	0.007*** (0.00)	0.006*** (0.00)	0.018*** (0.00)	0.028*** (0.00)	0.028*** (0.00)	0.009*** (0.00)	0.006*** (0.00)	0.005*** (0.00)	0.022*** (0.00)	0.016*** (0.00)	0.014*** (0.00)
Math Level	0.041* (0.02)	0.012 (0.01)	0.012 (0.01)	0.006 (0.02)	-0.012 (0.01)	-0.010 (0.01)	-0.008 (0.01)	-0.003 (0.01)	-0.004 (0.01)	0.013 (0.01)	0.035*** (0.01)	0.031*** (0.01)
Backward Induction	0.001 (0.02)	0.036 (0.02)	0.026 (0.02)	-0.009 (0.03)	0.011 (0.01)	0.009 (0.01)	0.014 (0.02)	-0.009 (0.02)	-0.008 (0.02)	-0.005 (0.02)	0.021 (0.02)	0.023 (0.01)
Financial Literacy	0.084** (0.03)	-0.009 (0.03)	-0.009 (0.02)	0.027 (0.03)	-0.020 (0.02)	-0.019 (0.02)	0.038 (0.03)	-0.004 (0.03)	-0.003 (0.03)	0.020 (0.03)	-0.009 (0.02)	-0.009 (0.02)
Control Questions	0.000 (0.01)	0.006 (0.01)	0.006 (0.01)	0.011 (0.01)	0.022*** (0.01)	0.021*** (0.01)	0.006 (0.01)	0.014* (0.01)	0.014 (0.01)	0.002 (0.02)	0.000 (0.01)	0.000 (0.01)
<i>Calculator Usage</i>												
Current Period		0.074*** (0.02)	0.008 (0.02)		0.004 (0.03)	0.003 (0.03)		-0.062** (0.02)	-0.075*** (0.03)		-0.001 (0.03)	-0.025 (0.03)
Partial		0.116*** (0.02)	0.050** (0.02)		0.048* (0.03)	0.043* (0.03)		0.000 (0.03)	-0.025 (0.03)		-0.021 (0.03)	-0.022 (0.03)
Sequential		0.028 (0.02)	-0.043* (0.02)		-0.049 (0.03)	-0.058* (0.03)		-0.001 (0.02)	-0.013 (0.02)		0.006 (0.03)	0.009 (0.03)
Complete		0.069*** (0.02)	-0.003 (0.02)		0.084*** (0.02)	0.075*** (0.02)		-0.039* (0.02)	-0.057** (0.02)		0.061** (0.03)	0.051* (0.03)
Constant	0.285 (0.42)	0.481 (0.47)	0.468 (0.45)	-0.159 (0.68)	-0.557 (0.36)	-0.517 (0.36)	0.388 (0.43)	0.340 (0.44)	0.357 (0.44)	0.473 (0.99)	0.447 (0.56)	0.440 (0.53)
Observations	1320	1320	1320	1290	1440	1440	1380	1440	1440	1440	1350	1350
Overall R^2	0.19	0.17	0.21	0.08	0.24	0.26	0.09	0.09	0.10	0.10	0.17	0.18

Notes: This table displays the outcomes of a random effects panel regression analysis. The regression models incorporate demographic control variables, including gender, age, prior experimental experience, academic degree, affiliated institution, current year of study, and risk tolerance levels. Additionally, we control for participants' comprehension of the instructions, by adding the participants' performance in the control questions. Standard errors are presented in parentheses. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

C Short-Span Planning Model's Estimation

Our methodology relies on the notion that an individual's initial consumption choice is influenced by their planning horizon and includes a random error component. This error component, denoted as ε , follows a normal distribution with a mean of zero and a standard deviation of σ . Therefore, the actual consumption choice for a specific planning horizon, denoted by c_H , is given by: $c_H = c_H + \varepsilon$, where H varies from 1 to 10 and c_H represents the theoretically optimal consumption level in Period 1 for the respective H . Additionally, we consider the possibility that a subset of subjects may make their consumption choices randomly. We model this by assuming that these subjects' decisions follow a uniform distribution ranging from 0 to 10. As a result, our model considers the population to comprise $H + 1$ distinct planner types, each characterized by a proportion p_h for h in the range of $[0,10]$.

To define a log-likelihood function, we integrate these proportions with the corresponding conditional densities. This methodology enables the computation of the sample log-likelihood for the n observations in our dataset. The log-likelihood function is expressed as follows:

$$\text{LogL} = \sum_{i=1}^n \ln \left[\frac{\hat{p}_0}{100} + \sum_{h=1}^{10} \hat{p}_h \frac{1}{\sigma} \phi \left(\frac{c_{i,R} - c_H^*}{\hat{\sigma}} \right) \right],$$

where $c_{i,R}$ represents the actual consumption of subject i in repetition R . The next step involves determining the proportions $\hat{p}_h$ and the standard deviation $\hat{\sigma}$ that maximize this log-likelihood function.

C.1 Individual-Level Analysis

To assess the model's predictions against actual planning horizons on an individual basis, we employ $\hat{\sigma}$ to establish a range for c_H for each participant, thereby facilitating the assignment of corresponding planning horizons. Figure C.1 illustrates this detailed examination. The x -axis shows the horizon predicted by the short-span planning model,

Table C.1: Optimal Consumption Level in Period 1, c_H

H	ConstantYP	PosIR	FluctY	FluctP
1	10.00	10.00	15.00	8.00
2	10.00	9.09	10.00	4.00
3	10.00	8.26	11.67	4.80
4	10.00	7.51	10.00	4.00
5	10.00	6.83	11.00	4.44
6	10.00	6.21	10.00	4.00
7	10.00	5.64	10.71	4.31
8	10.00	5.13	10.00	4.00
9	10.00	4.67	10.56	4.24
10	10.00	4.24	10.00	4.00

Table C.2: Estimation Outcome

Parameter	Coefficient	Standard Error
$\hat{p}_1$	0.21***	0.034
$\hat{p}_2$	0.118***	0.028
$\hat{p}_3$	0.071***	0.022
$\hat{p}_4$	0.035**	0.016
$\hat{p}_5$	0.032**	0.015
$\hat{p}_6$	0.025*	0.014
$\hat{p}_7$	0	NA
$\hat{p}_8$	0.116***	0.027
$\tilde{p}_9$	0.001	0.021
$\hat{p}_{10}$	0.099***	0.031
Uncategorized	0.291***	0.04
$\hat{\sigma}$	0.162***	0.013

Note: p_7 was set to 0 because otherwise the function did not converge. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

whereas the y -axis presents the mode of the horizons used by subjects with the budgeting calculator during the first period of each repetition. Notably, the patterns observed herein provide further insight: the actual planning horizons are typically more extended than those predicted by the short-term horizon model. Moreover, in over 25 percent of cases, the model either fails to categorize an individual or predicts a horizon of 1, while subjects opt for the full horizon. Conversely, subjects adhere to the model's full-horizon prediction in only 3 percent of cases. Nevertheless, in approximately 3 percent of instances, the model forecasts full horizon usage, yet individuals select a shorter horizon.




Figure C.1: Comparison between the predicted and the observed horizon at the individual level

D Hand-to-Mouth Heuristic

One evident heuristic that subjects may adopt to circumvent complexity is the "hand-to-mouth behavior (H2M)," where they expend all their available cash each period. Yet, this approach comes with the cost of potential utility loss. When subjects refrain from saving in any period, they incur a utility reduction of 3.6% in PosIR and 3.4% in both FluctP and FluctY, relative to the optimal path.²³ In the ConstantYP treatment, the optimal consumption path inherently involves zero savings per period. However, even in this treatment, we can examine the adoption of this heuristic by observing deviations from the unconditionally optimal consumption in at least one period. Two relevant questions arise: Does this heuristic become more prevalent as complexity increases, and does access to the calculator limit its usage?

To address these questions, Figure D.1 displays the percentage of subjects using the H2M heuristic across various treatments and over time. In the ConstantYP, FluctY, and FluctP treatments, subjects who chose not to save, given that not saving was the implied optimal choice for that specific period, are excluded from the analysis. The figure demonstrates that the adoption of this heuristic varies across treatments.

In the ConstantYP treatment, the percentage of subjects employing the H2M heuristic is small, accounting for less than 10 percent of the periods. Furthermore, there is no discernible distinction between subjects with calculator access and those without in this treatment.

In the FluctY treatment, subjects display the correct intuition by employing the H2M heuristic when their income is low, aligning with consumption smoothing principles. Here, the heuristic was utilized in 21.99 percent of the periods when the calculator was present and in 19.08 percent of the periods when it was absent.²⁴

In the context of complex environments, starting with the PosIR treatment, the absence of a calculator led to a higher incidence of the Hand to Mouth (H2M) heuristic.

²³In 9% of the repetitions, participants used this heuristic in all periods.

²⁴Running a test on the equality of proportions at a 95% significance level, we do not find a statistically significant difference between the two calculator conditions ($p = 0.0699$).

The H2M heuristic was utilized in 12.03 percent of the periods when the calculator was available, and in 16.02 percent of the periods when it was not (across all repetitions). This pattern underscores the potential utility of the calculator in mitigating reliance on H2M behavior. Notably, there were 6 subjects who consistently used this heuristic from the beginning to the end of Repetition 3 when the calculator was not enabled, which is double the number observed when the calculator was available.

In the FluctP treatment, the H2M heuristic was employed in 22.46 percent of the periods with calculator access, and in 17.05 percent of periods without it, across all repetitions.²⁵ Notably, within this treatment, the proportion of subjects utilizing the heuristic increases during periods of low prices, reflecting their awareness of the benefits of purchasing more units at lower costs, though they often exceed optimal consumption levels. Furthermore, considering participants who consistently applied the heuristic throughout Repetition 3, only three subjects did so with the calculator enabled, compared to seven without calculator access.

²⁵We reject the null hypothesis of equality of proportions ($p = 0.0006$). A plausible explanation for this behavior in the FluctP treatment is that employing the H2M heuristic requires more complex arithmetic operations—specifically, dividing numbers by 50 or 150—compared to other treatments, which typically involve divisions by 100. The budgeting calculator may encourage the adoption of the H2M heuristic by simplifying these arithmetic challenges.




Figure D.1: Percentage of Subjects Using the Hand-to-Mouth Heuristic

E The Effect of the Calculator on Payoffs

In this appendix, we illustrate the effect of using a calculator on payoffs. We also use random choices as a secondary benchmark for comparison.

The table below shows the expected payoffs in dollar terms for treatments with and without a calculator, as well as for random consumption choice simulations. We refer to these as expected payoffs because, for each subject, we averaged the payoffs over three repetitions. This is because subjects were only paid based on one randomly selected repetition at the end of the experiment.

For the random choice simulations, we conducted 1,000,000 simulations for each treatment environment and repetition. We then calculated the payoffs and took the average over the repetitions. In these simulations, we ensured that borrowing constraints were binding and that the cash balance was exhausted in consumption during the last period.

From the table, we observe the following that subjects with access to a calculator performed better than those without in the ConstantYP and PosIR treatments, although the difference is minimal in FluctY and FluctP. Interestingly, in the PosIR treatment, subjects with and without a calculator performed worse than the average simulation for the random choice model. This difference is more pronounced in the No Calculator sessions.

Treatment	Average payoffs in CAD (\$)			No Calc vs. Random	Calc vs. Random	Calc vs. No Calc
	No Calculator	Calculator	Random Choice	(% Chg Norm by by Max Payoffs)	(% Chg Norm by by Max Payoffs)	(% Chg Norm by by Max Payoffs)
ConstantYP	8.87	9.2	8.95	-0.83%	2.59%	3.42%
PosIR	8.66	9.1	9.28	-6.19%	-1.80%	4.39%
FluctY	9.19	9.12	8.95	2.49%	1.76%	-0.72%
FluctP	8.89	8.95	8.81	0.83%	1.45%	0.62%

F Individual Time Series




Figure F.1: Constant YP No Calc, Repetition 1 (Session 1)




Figure F.2: Constant YP No Calc, Repetition 3 (Session 1)




Figure F.3: Constant YP No Calc, Repetition 1 (Session 2)




Figure F.4: Constant YP No Calc, Repetition 3 (Session 2)




Figure F.5: Constant YP Calc, Repetition 1 (Session 1)




Figure F.6: Constant YP Calc, Repetition 3 (Session 1)




Figure F.7: Constant YP No Calc, Repetition 1 (Session 2)




Figure F.8: Constant YP Calc, Repetition 3 (Session 2)




Figure F.9: Pos IR No Calc, Repetition 1 (Session 1)




Figure F.10: Pos IR No Calc, Repetition 3 (Session 1)




Figure F.11: Pos IR No Calc, Repetition 1 (Session 2)




Figure F.12: Pos IR No Calc, Repetition 3 (Session 2)




Figure F.13: Pos IR Calc, Repetition 1 (Session 1)




Figure F.14: Pos IR Calc, Repetition 3 (Session 1)




Figure F.15: Pos IR No Calc, Repetition 1 (Session 2)




Figure F.16: Pos IR Calc, Repetition 3 (Session 2)




Figure F.17: FluctY No Calc, Repetition 1 (Session 1)




Figure F.18: FluctY No Calc, Repetition 3 (Session 1)




Figure F.19: FluctY No Calc, Repetition 1 (Session 2)




Figure F.20: FluctY No Calc, Repetition 3 (Session 2)




Figure F.21: FluctY Calc, Repetition 1 (Session 1)




Figure F.22: FluctY Calc, Repetition 3 (Session 1)




Figure F.23: FluctY No Calc, Repetition 1 (Session 2)



Figure F.24: FluctY Calc, Repetition 3 (Session 2)




Figure F.25: FluctP No Calc, Repetition 1 (Session 1)




Figure F.26: FluctP No Calc, Repetition 3 (Session 1)




Figure F.27: FluctP No Calc, Repetition 1 (Session 2)




Figure F.28: FluctP No Calc, Repetition 3 (Session 2)




Figure F.29: FluctP Calc, Repetition 1 (Session 1)




Figure F.30: FluctP Calc, Repetition 3 (Session 1)




Figure F.31: FluctP No Calc, Repetition 1 (Session 2)




Figure F.32: FluctP Calc, Repetition 3 (Session 2)

G Instructions

G.1 Main Instructions

G.1.1 Calc Sessions

The instructions distributed to subjects in treatments where the calculator was enabled are reproduced on the subsequent pages. All subjects received identical instructions, with the exception of those in the PosIR treatment, where minor adjustments were made to incorporate interest rates.

You are taking part in an economics experiment in which you will be able to earn money. Your earnings will depend on your decisions. It is therefore important to read these instructions with attention. During the experiment you are not allowed to communicate with any other participant. If you do not follow these instructions you will be excluded from the experiment and receive only the show-up payment of \$7.

The experiment consist of **4 PARTS**. These instructions are for the first part only. After completing each part of the experiment, you will receive instructions for each subsequent part. The earnings you accumulate will be added up at the end of the experiment, and converted to Canadian dollars. Specifically, for every **25 points** you accumulate, you will obtain **\$1**. You will also receive a \$7 show-up payment (if you arrived on time). Before you leave the lab, you will sign a receipt and will be paid in cash privately.

First Part

In the first part of this experiment, you will be making decisions on how much to save and spend over a number of periods.

There are two objects of interest in this experiment, **tokens** and **points**. The total number of points you accumulate in a repetition will determine your monetary payoff.

You will participate in 3 repetitions of the exact same experiment, each consisting of 10 periods. In every repetition, at the beginning of each period, you will receive some tokens. In addition to these tokens, you may have additional tokens saved from previous periods.

Purchase Decision

Each period, after viewing the total number of tokens you have available, you must decide how many units of output you would like to buy using tokens. When you submit your purchases orders you can use up to two decimal places (the minimum you can buy is 0). Output is sold at a certain price per unit. The output you buy will be transformed into points. The more output you acquire in a period, the higher will be your points earnings for that period. Importantly, as you purchase more output in a single period, you will earn fewer and fewer additional points. Your first unit will be worth the most, and each subsequent point will be worth less.

After submitting your purchase order, the computer will calculate your expenditure as following:

$$\text{Expenditure} = \text{Number of units purchased} \times \text{Price per unit.}$$

This expenditure will be deducted from your token balance. If, at a certain period, you do not have enough tokens to buy output, you will not be able to complete your order. You may not spend more than your token balance.


Your token balance at the start of each period is given by the following:

$$\text{Tokens at the beginning of current period} = \text{Tokens from previous period} + \text{Current income}$$

Compensation

The points will be converted into Canadian dollars at the end of the experiment. You will be compensated according to the following rules:

1. The game will be repeated 3 times. At the end of the experiment, the computer will randomly select one of the repetitions for payment. That is, there is an equal chance that any repetition will be the one that counts for payment.
2. The diagram below shows the relation between purchased output, points, and cash (\$).



3. The amount of points you earn in the randomly selected repetition will be converted into CAD at the rate:

$$\mathbf{25 \text{ points} = \$1}$$

4. Any tokens held at the end of a repetition are worthless.
5. Additionally, you will receive 4 points for every control question you answer correctly in the first attempt; 3 points for every question you answer correctly in the second attempt; and 2 points for every question you answer correctly in the third attempt.

Information

You will be provided precise information about the income and prices you will face in all 10 periods of each repetition. That is, there is no uncertainty about these variables. You will also have information about your previous decisions and economic variables. Please note that prices and income may vary from repetition to repetition.

Repetitions

The experiment will consist of 3 repetitions of 10 periods each. After a repetition is completed, you will see a Review Screen that will display your total points from that repetition. There will be no carryover of tokens or points between repetitions. When a new repetition begins, all token balances and points will be reset to zero.

Output to Points Converter

Throughout the experiment, you will have access to an Output to Points Converter that you may use to help you make decisions. To use the Output to Points Converter, you will need to enter the number of units of output you wish to convert to points. After clicking the OK button, the computer will display the points associated with the output you entered.

Output to Points Converter
Enter the units of output you would like to convert into points and click OK.

Standard Calculator

You may use a standard calculator by clicking on



Payoff Calculator

Throughout the experiment, you will have access to a calculator that you may use to help you make decisions.

To use the calculator, you will need to fill in hypothetical values for your purchase decisions in the current and future periods. After all your hypothetical decisions have been submitted, you will be able to see what your points and tokens balance would be. You can consider as many hypothetical combinations as you want before making each decision.

Before the experiment starts you will learn how to use the calculator; you will be able to practice with it; and finally, you will have to answer some paid control questions.

Decision Box
Please submit your choice for Period 1

Number of units you would like to purchase

Instructions
Press the button below to open the Instructions

Calculator

Period	Income (Tokens)	Price per Unit	Token Balance (beginning of the period)	Purchased Units	Points	Accumulated Points	Final Tokens Balance
1							
2							
3							
4							
5							
6							
7							
8							
9							
10							

Standard Calculator

Remember that your actual purchase decision has to be entered on the left hand side of the screen.

G.1.2 No Calc Sessions

For the sessions in which the budgeting calculator was not enabled the third page of the instructions was adapted as shown below.

Output to Points Converter

Throughout the experiment, you will have access to an Output to Points Converter that you may use to help you make decisions. To use the Output to Points Converter, you will need to enter the number of units of output you wish to convert to points. After clicking the OK button, the computer will display the points associated with the output you entered.

Output to Points Converter
Enter the units of output you would like to convert into points and click OK.

Standard Calculator

You may use a standard calculator by clicking on



Income and Expenditure Table

Throughout the experiment, you will have access to an Income and Expenditure Table that you may use to help you make decisions.

Each period, you will be able to see what your points and tokens balance are.

Before the experiment starts, you will learn how to read the table. You will also have to answer some paid control questions in which you will be able to put to the test your ability to read the table.

Decision Box
Please submit your choice for Period 1

Number of units you would like to purchase

Instructions
Press the button below to open the Instructions

Income and Expenditure Table

Period	Income (Tokens)	Price per Unit	Token Balance (beginning of the period)	Purchased Units	Points	Accumulated Points	Final Tokens Balance
1							
2							
3							
4							
5							
6							
7							
8							
9							
10							

Standard Calculator

Remember that your actual purchase decision has to be entered on the left hand side of the screen.

G.2 Interactive Instructions




Figure G.1: Screenshot for Interactive Instructions 1




Figure G.2: Screenshot for Interactive Instructions 2




Figure G.3: Screenshot for Interactive Instructions 3




Figure G.4: Screenshot for Interactive Instructions 4

Interactive Instructions (2/2)

How to use the Calculator

-- The 3rd column shows the price per unit of output for each period.

Back

Next

Calculator

Period	Income (Tokens)	Price per Unit
1	1000.00	100.00
2	1000.00	100.00
3	1000.00	100.00
4	1000.00	100.00
5	1000.00	100.00
6	1000.00	100.00
7	1000.00	100.00
8	1000.00	100.00
9	1000.00	100.00
10	1000.00	100.00

Figure G.5: Screenshot for Interactive Instructions 5

Interactive Instructions (2/2)

How to use the Calculator

-- The 4th column shows the interest rate you will receive for your savings.
-- This rate will not change throughout the experiment. It will always be 10%.

Back

Next

Calculator

Period	Income (Tokens)	Price per Unit	Interest Rate (%)
1	1000.00	100.00	10.00
2	1000.00	100.00	10.00
3	1000.00	100.00	10.00
4	1000.00	100.00	10.00
5	1000.00	100.00	10.00
6	1000.00	100.00	10.00
7	1000.00	100.00	10.00
8	1000.00	100.00	10.00
9	1000.00	100.00	10.00
10	1000.00	100.00	10.00

Figure G.6: Screenshot for Interactive Instructions 6

Interactive Instructions (2/2)

How to use the Calculator

-- The 5th column shows your Token Balance at the beginning of the period, in other words it show you how many tokens you may use for spending. To calculate it, the computer adds your current income, your Token Balance in the previous period, and your interest income.
-- Note that in period 1 this amount is equal to your income. This is because you do not have any savings from previous periods.

Back

Next

Calculator

Period	Income (Tokens)	Price per Unit	Interest Rate (%)	Token Balance (beginning of the period)
1	1000.00	100.00	10.00	1000.00
2	1000.00	100.00	10.00	
3	1000.00	100.00	10.00	
4	1000.00	100.00	10.00	
5	1000.00	100.00	10.00	
6	1000.00	100.00	10.00	
7	1000.00	100.00	10.00	
8	1000.00	100.00	10.00	
9	1000.00	100.00	10.00	
10	1000.00	100.00	10.00	

Figure G.7: Screenshot for Interactive Instructions 7

Interactive Instructions (2/2)

How to use the Calculator

-- In the 6th column you will be able to enter the number of units of output you wish to buy in each of the periods.
-- Every period, the maximum amount of units you can purchase will depend on your Token Balance and the Price per unit. You may not spend more than your Token Balance.

Back

Next

Calculator

Period	Income (Tokens)	Price per Unit	Interest Rate (%)	Token Balance (beginning of the period)	Purchased Units
1	1000.00	100.00	10.00	1000.00	
2	1000.00	100.00	10.00		
3	1000.00	100.00	10.00		
4	1000.00	100.00	10.00		
5	1000.00	100.00	10.00		
6	1000.00	100.00	10.00		
7	1000.00	100.00	10.00		
8	1000.00	100.00	10.00		
9	1000.00	100.00	10.00		
10	1000.00	100.00	10.00		

Figure G.8: Screenshot for Interactive Instructions 8

Interactive Instructions (2/2)

How to use the Calculator

-- The 7th column displays the number of points you obtain for your output's purchases.
-- Importantly, as you purchase more output in a single period, you will earn fewer and fewer points.

Back

Next

Calculator

Period	Income (Tokens)	Price per Unit	Interest Rate (%)	Token Balance (beginning of the period)	Purchased Units	Points
1	1000.00	100.00	10.00	1000.00		
2	1000.00	100.00	10.00			
3	1000.00	100.00	10.00			
4	1000.00	100.00	10.00			
5	1000.00	100.00	10.00			
6	1000.00	100.00	10.00			
7	1000.00	100.00	10.00			
8	1000.00	100.00	10.00			
9	1000.00	100.00	10.00			
10	1000.00	100.00	10.00			

Figure G.9: Screenshot for Interactive Instructions 9

Interactive Instructions (2/2)

How to use the Calculator

-- The 8th column shows you the number of points you have accumulated from period 1 up to the current period.
-- Remember that at the end of the experiment you will be compensated based on the points you obtain in a randomly selected repetition.

Back

Next

Calculator

Period	Income (Tokens)	Price per Unit	Interest Rate (%)	Token Balance (beginning of the period)	Purchased Units	Points	Accumulated Points
1	1000.00	100.00	10.00	1000.00			
2	1000.00	100.00	10.00				
3	1000.00	100.00	10.00				
4	1000.00	100.00	10.00				
5	1000.00	100.00	10.00				
6	1000.00	100.00	10.00				
7	1000.00	100.00	10.00				
8	1000.00	100.00	10.00				
9	1000.00	100.00	10.00				
10	1000.00	100.00	10.00				

Figure G.10: Screenshot for Interactive Instructions 10

91

Interactive Instructions (2/2)

How to use the Calculator

-- The 9th column shows you balance in tokens at the end of each period.
-- The computer calculates it by subtracting your expenditure (no. of units of output you purchased X price per unit) from you initial Token Balance.

Back

Next

Calculator

Period	Income (Tokens)	Price per Unit	Interest Rate (%)	Token Balance (beginning of the period)	Purchased Units	Points	Accumulated Points	Final Tokens Balance
1	1000.00	100.00	10.00	1000.00				
2	1000.00	100.00	10.00					
3	1000.00	100.00	10.00					
4	1000.00	100.00	10.00					
5	1000.00	100.00	10.00					
6	1000.00	100.00	10.00					
7	1000.00	100.00	10.00					
8	1000.00	100.00	10.00					
9	1000.00	100.00	10.00					
10	1000.00	100.00	10.00					

Figure G.11: Screenshot for Interactive Instructions 11

Interactive Instructions (2/2)

Please take some time to practice with the calculator.

Go to Control Questions

Calculator

Period	Income (Tokens)	Price per Unit	Interest Rate (%)	Token Balance (beginning of the period)	Purchased Units	Points	Accumulated Points	Final Tokens Balance
1	1000.00	100.00	10.00	1000.00	2	9.49	9.49	800.00
2	1000.00	100.00	10.00	1880.00	3	11.62	21.11	1580.00
3	1000.00	100.00	10.00	2738.00	13	24.19	45.29	1438.00
4	1000.00	100.00	10.00	2582.00	13	24.19	69.48	1282.00
5	1000.00	100.00	10.00	2410.00	18	25.98	95.46	910.00
6	1000.00	100.00	10.00	2001.00	20	30.00	125.46	1.00
7	1000.00	100.00	10.00	1001.00	2	9.49	134.95	801.00
8	1000.00	100.00	10.00	1881.00				
9	1000.00	100.00	10.00					
10	1000.00	100.00	10.00					

Clear Values

Calculate

Standard Calculator

Figure G.12: Screenshot for Interactive Instructions 12

92

H Computer Interface

H.1 Control Questions

Control Questions (1/4)

1. How many times the game will be repeated? 3 (Correct Answer)

2. Of how many periods consists each repetition? 10 (Correct Answer)

Next Questions

Instructions
Press the button below to open the instructions
Open Instructions

Figure H.13: Screenshot for Control Question 13

Control Questions (2/4)

3. How many would you obtain if you purchased 1 unit of output in a single period? 6.71 (Correct Answer)

4. How many would you obtain if you purchased 2 units of output in a single period? 9.49 (Correct Answer)

5. How many would you obtain if you purchased 9 units of output in a single period? 20.12 (Correct Answer)

6. How many would you obtain if you purchased 10 units of output in a single period? 21.21 (Correct Answer)

7. Your points increase more when (select one of the options below):

☒ A. You increase your output purchases from 1 to 2 units.

☐ B. You increase your output purchases from 9 to 10 units.

☐ C. The increase is the same in both cases.

Submit

Output to Points Converter
Enter the units of output you would like to convert into points and click OK.

OK

If you purchase 10.00 units you will obtain 21.21 Points.

Purchased Units	Points
1.00	6.71
2.00	9.49
9.00	20.12
10.00	21.21

Instructions
Press the button below to open the instructions
Open Instructions

Figure H.14: Screenshot for Control Question 2

Control Questions (3/4)

8. Suppose it is period 1. What is the maximum number of units of output you can purchase in this period? 10.00 (Correct Answer)

9. What is the minimum number of units you can purchase in period 1? 0.00 (Correct Answer)

10. Suppose it is period 5. Suppose also that you have submitted the purchase decisions shown in the table below. What is the maximum amount of output you will be able to convert into points in period 5?

☐ A. 10.00 units of output
☐ B. 0 units of output
☒ C. 30.04 units of output

Submit

Calculator

Period	Income (Tokens)	Price per Unit	Interest Rate (%)	Token Balance (beginning of the period)	Purchased Units	Points	Accumulated Points	Final Tokens Balance
1	1000.00	100.00	10.00	1000.00	8.00	18.97	18.97	209.00
2	1000.00	100.00	10.00	1220.00	4.00	13.42	13.42	826.00
3	1000.00	100.00	10.00	1962.00	7.00	17.75	17.75	1262.00
4	1000.00	100.00	10.00	2322.00	5.00	15.00	15.00	1822.00
5	1000.00	100.00	10.00	3004.00				
6	1000.00	100.00	10.00					
7	1000.00	100.00	10.00					
8	1000.00	100.00	10.00					
9	1000.00	100.00	10.00					
10	1000.00	100.00	10.00					

Clear Values
Calculate
Standard Calculator

Instructions

Press the button below to open the instructions

[Open Instructions](#)

Figure H.15: Screenshot for Control Question 3

Control Questions (4/4)

11. Suppose it is period 10. Suppose also that you have submitted the purchase decisions shown in the table below. What is the maximum amount of output you will be able to convert into points in period 10? 95.24 (Correct Answer)

12. If you choose to spend less than 9524.00 tokens in period 10.

☐ A. You will be able to carry over your savings
☐ B. Your savings will have some value at the end of the experiment
☒ C. Your savings will be worthless at the end of the experiment

Submit

Calculator

Period	Income (Tokens)	Price per Unit	Interest Rate (%)	Token Balance (beginning of the period)	Purchased Units	Points	Accumulated Points	Final Tokens Balance
1	1000.00	100.00	10.00	1000.00	3.00	11.62	11.62	700.00
2	1000.00	100.00	10.00	1770.00	4.00	13.42	13.42	1370.00
3	1000.00	100.00	10.00	2567.00	6.00	16.43	16.43	1967.00
4	1000.00	100.00	10.00	3098.00	4.00	13.42	13.42	2698.00
5	1000.00	100.00	10.00	3968.00	8.00	16.43	16.43	3368.00
6	1000.00	100.00	10.00	4705.00	4.00	13.42	13.42	4305.00
7	1000.00	100.00	10.00	5738.00	2.00	9.49	9.49	5538.00
8	1000.00	100.00	10.00	7090.00	5.00	15.00	15.00	6590.00
9	1000.00	100.00	10.00	8249.00	5.00	15.00	15.00	7749.00
10	1000.00	100.00	10.00	9524.00	95.24	65.47	65.47	0.00

Clear Values
Calculate
Standard Calculator

Instructions

Press the button below to open the instructions

[Open Instructions](#)

Figure H.16: Screenshot for Control Question 4

H.2 Risk Preferences Elicitation

Part 2

In this part of the experiment you will select from among six different gambles the one gamble you would like to play. The six different gambles are listed on the table below. You must select one and only one of these gambles. Each gamble has two possible outcomes (Event A or Event B) with the indicated probabilities of occurring. Your compensation for this part of the study will be determined by: 1) which of the six gambles you select; and 2) which of the two possible events occur.

For example: If you select gamble 4 and Event A occurs, you will earn 16 points. If Event B occurs, you will earn 52 points.

For every gamble, each event has a 50% chance of occurring.

After you have selected your gamble you will roll a six-sided virtual dice to determine which event will occur. If you roll a 1, 2, or 3, Event A will occur. If you roll a 4, 5, or 6, Event B will occur.

Gamble	Event	Payoff (Points)	The event occurs if you roll	Probabilities
1	A	28	1, 2, or 3	50%
	B	28	4, 5, or 6	50%
2	A	24	1, 2, or 3	50%
	B	36	4, 5, or 6	50%
3	A	20	1, 2, or 3	50%
	B	44	4, 5, or 6	50%
4	A	16	1, 2, or 3	50%
	B	52	4, 5, or 6	50%
5	A	12	1, 2, or 3	50%
	B	60	4, 5, or 6	50%
6	A	2	1, 2, or 3	50%
	B	70	4, 5, or 6	50%

Submit

Standard Calculator



Figure H.17: Screenshot for Stage 2

H.3 Financial Literacy

Part 3

1. Suppose you had \$100 in a savings account and the interest rate was 2% per year. After 5 years, how much do you think you would have in the account if you left the money to grow?

☐ A. More than \$102
☐ B. Exactly \$102
☐ C. Less than \$102

Standard Calculator 

2. Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After 1 year, would you be able to buy more than, exactly the same as, or less than today with the money in this account?

☐ A. More than today
☐ B. Exactly the same as today
☐ C. Less than today

Standard Calculator 

3. Do you think that the following statement is true or false? "Buying a single company stock usually provides a safer return than a stock mutual fund."

☐ A. True
☐ B. False
☐ C. There is not enough information

Standard Calculator 

4. Do you think the following statement is true or false: A 15-year mortgage typically requires higher monthly payments than a 30-year mortgage, but the total interest over the life of the loan will be less?

☐ A. True
☐ B. False
☐ C. There is not enough information

Standard Calculator 

5. If interest rates rise, what will typically happen to bond prices?

☐ A. They will rise
☐ B. They will fall
☐ C. They will stay the same
☐ D. There is no relationship

Standard Calculator 

Figure H.18: Screenshot for Stage 3

H.4 Race to 60

Part 4

Instructions

In this game, you play 8 repetitions of the game "Race to 60". Your goal is to win against the computer. In the game, you and the computer alternately choose numbers between 1 and 10. The numbers are added up, and whoever chooses the number that pushes the sum of the numbers to or above 60, wins the game.

Specifically, at the beginning of each game, you choose a number between 1 and 10 (both included). Then the game follows these steps: The computer enters a number between 1 and 10. This number is added to your number. The sum of all chosen numbers up the current round is shown on the screen. If the sum is smaller than 60, you enter a number between 1 and 10, which in turn will be added to all number chosen up to the current round by you and the computer. This sequence is repeated until the sum of all numbers is greater or equal than 60. Whoever (i.e. you or the computer) chooses the number that adds up to a sum equal or above 60, wins the game.

You will be playing this game 8 times. For each game won, you receive 8 points.

Please enter a number between 1 and 10 and press Submit.

Submit

Your choice is: 4

The computer's choice is: 6

The sum of the numbers entered so far is:

32

History of Choices

Round	Your Choice	Computer's Choice	Sum
1	1	8	9
2	3	10	22
3	4	6	32

Standard Calculator 

Figure H.19: Screenshot for Stage 3