



Teacher guide: public debt, consumption and generations

Laboratory implementation, theoretical foundations,
experimental evidence, and variants

Behavioral and Experimental Economics Research Lab
Department of Economics



Revised edition | October 2026

<https://beer-lab.org/>

How to use this guide

This guide supports the preparation, implementation, and classroom discussion of the family of consumption, public-debt and transfer activities available through the Behavioral and Experimental Economics Research Lab (BEER Lab) at the Department of Economics of Tecnológico de Monterrey. It is intended for instructors with different levels of familiarity with experimental and behavioral economics.

The document offers three reading paths. The **quick guide** is enough to launch a session. The operational chapters explain each screen used by instructors and students. The conceptual chapters present the theories, predictions, and empirical findings needed to lead the post-game discussion.

Access instructions and common screens are explained once. Each variant then focuses only on what changes relative to the basic game: configuration, group size, decisions, predictions, relevant outcomes, and academic literature.

Author

This guide was prepared by Stanislaw Maldonado, research professor in the Department of Economics at Tecnológico de Monterrey, Monterrey Campus. For questions or suggestions about its content, please contact the author at stanislao.maldonado@tec.mx.

For the instructor

This guide follows the English-language interface. Spanish configurations use the same logic. Never project individual codes or connect outcomes to names, student identification numbers, or seats in the room.

Contents

How to use this guide	i
1 Preparation and teaching purpose	1
1.1 Access through the updated laboratory website	1
1.2 The question behind the activity	5
1.3 Product facts	5
1.4 Observable learning objectives	6
1.5 Preparation before students arrive	6
1.6 A minimum classroom sequence	7
2 Baseline budgets and predictions	8
2.1 Two dates and an identical public purchase	8
2.2 An example without calculus	8
2.3 Present value and available wealth	9
2.4 Preferences and the benchmark solution	9
2.5 What the activity establishes	10
3 Access, sessions and basic modes	11
3.1 From the portal to the correct room	11
3.2 Basic configurations and choice order	11
3.3 Create a session and distribute links	12
3.4 Local testing and product routes	15
4 Participant screens and actions	16
4.1 Welcome and comprehension	16
4.2 The individual decision screen	18
4.3 Outcomes and transition	20
4.4 Roles and intergenerational waiting	21
4.5 Explicit completion	23
5 Extensions: credit, risk and information	24
5.1 Compare one mechanism at a time	24
5.2 No private borrowing	24
5.3 Different borrowing and deposit rates	25
5.4 Uncertain future tax	27
5.5 Present-value accounting support	27
5.6 Temporal weight and zero interest	30
6 Transfers between generations and research evidence	31
6.1 One pair and two consumption recipients	31
6.2 Full weight and an operative bequest	31
6.3 Zero and intermediate weights	31

6.4	Fixed pairs and between-block rematching	33
6.5	Relationship to research	33
7	Reading the report and leading discussion	34
7.1	Private access and controlled revelation	34
7.2	Calendar comparison and complete blocks	34
7.3	Benchmarks, distances and utility	38
7.4	Distribution and evolution	38
7.5	External accounts and recipient outcomes	40
8	Lesson plans, exercises and data use	41
8.1	A 45-minute introductory lesson	41
8.2	A 75-minute intermediate lesson	41
8.3	Exercises with answer guidance	42
8.3.1	An apparently larger obligation	42
8.3.2	Negative saving and final consumption	42
8.3.3	Changing the temporal weight	42
8.3.4	A fund in two states	42
8.3.5	Partial compensation through transfers	42
8.3.6	Two means with an incomplete sample	43
8.3.7	Interest and the income profile	43
8.4	Data, privacy and inference limits	43
8.5	Common problems and an operating checklist	43

List of Figures

1.1	Current institutional website. Select CLASSROOM in the top navigation menu to open the teaching experiments. Captured on October 7, 2026.	1
1.2	Current catalog: search by title, topic, or JEL code. Then locate this family's card and select language and edition. Captured on October 7, 2026.	2
1.3	Demo contains only the configurations of this game family, in both languages and editions. Continuation 1/2.	3
1.5	Rooms contains the four rooms for this family: basic and extended editions in Spanish and English.	4
1.6	Session-creation form for the basic English room. Check the selected configuration and participant requirement before creating the session.	4
3.1	Select the English mode and confirm four test participants. The same procedure supports the actual class size.	13
3.2	The session panel provides participant entry and private session navigation. Test links use impersonal labels in this capture.	14
4.1	The welcome obtains acceptance and explains the task. Participants keep their own session links private.	16
4.2	Current instructions for the basic edition: initial resources, decision rules, and experimental units. Continuation 1/2.	17
4.4	Current comprehension check with an incorrect answer, its specific explanation, and retained responses.	18
4.5	Read income, taxes and interest before entering consumption. The announced interval is part of the budget rules.	19
4.6	Budget quantities and the performance score appear separately. Reconstruct final consumption from the displayed balance.	20
4.7	The role screen identifies the transfer decision-maker and the announced recipient weight. Check who pays the final tax.	21
4.8	The second generation forecasts an initial transfer without a reward. The response describes an expectation and does not authorise a different allocation.	22
4.9	Waiting synchronises only the pair. Delay indicates a missing response and reveals no economic choice.	22
4.10	Completion lists both calendars and accumulated points. Compare consumption and balances within a block before comparing different modes.	23
5.1	Consumption is limited by current resources after tax. Compare the upper bound with the unconstrained benchmark of 89.52 before interpreting choices.	25
5.2	Both rates are announced before choosing. A negative balance uses the borrowing rate and a positive balance uses the deposit rate.	26
5.3	The future-tax screen states both obligations, probabilities and the fiscal fund's role. The draw is disclosed after the choice.	28

5.4	The extra box displays resources in present-value terms. It organises existing information without adding income or changing taxes. Continuation 1/2.	29
6.1	The first generation chooses consumption from current funds and leaves a nonnegative transfer. The announced weight identifies the scoring rule.	32
7.1	Before completion, progress is visible and revelation is disabled. One finished decision is not a complete comparable block.	35
7.2	Completion enables the button but keeps outcomes hidden until the instructor reveals them.	36
7.3	Read calendar counts, matched-block count and within-block differences together. The denominator determines the comparison represented by a mean change.	37
7.4	The histogram displays concentration and heterogeneity. Counts represent decisions, rather than unique participants in a multi-episode session.	39
7.5	The table accompanies the curve with counts and episode means. Inspect calendar composition before interpreting any position.	39
7.6	Public accounts are separated from recipient points and forecast errors. Recipient outcomes state their corresponding record count.	40

Preparation and teaching purpose

1.1 Access through the updated laboratory website

Start at <https://beer-lab.org/en/>. Select **CLASSROOM** in the top navigation menu; it opens <https://teaching.beer-lab.org/en/>. In the **Classroom experiments** catalog, use the search box to locate the activity by name. You can also filter by **Topic** or **JEL code**. Clear any filters left from an earlier search before looking for another game.

Each family's card distinguishes **Español** and **English** access links. The **Show extended version** checkbox selects that edition before following a link. Catalog language and session language are separate choices: verify both. Students use the room specified by the instructor; the instructor prepares and monitors the session through this family's own administrative pages.




Figure 1.1: Current institutional website. Select CLASSROOM in the top navigation menu to open the teaching experiments. Captured on October 7, 2026.




Figure 1.2: Current catalog: search by title, topic, or JEL code. Then locate this family's card and select language and edition. Captured on October 7, 2026.

Checking resources, comprehension, and participant requirements

Before opening the room, check the selected configuration and the number of students present. The session-creation form states the group size and minimum required by that configuration. An individual activity accepts any positive participant count; an interactive activity must satisfy its stated requirement: complete groups when size is fixed, or a minimum participant count when size is flexible. Requirements may change across treatments, for example when adding a third party or changing an institution's size. If you edit a grouping parameter, check the requirement for the edited configuration, rather than relying only on the demo's suggested count.

In the instructions, ensure that students can identify what each role receives, what they can decide, and how the outcome is calculated. An initial account, a distributable budget, and final earnings need not be the same amount. Refer to the resources panel and the example for the selected treatment. Do not carry the basic edition's endowment into a treatment that changes accounts or introduces additional positions.

ECU stands for *Experimental Currency Units*: the activity's unit of account. An outcome of 100 ECU does not automatically equal 100 pesos or 100 dollars. Conversion depends on the session configuration. Before announcing a payment, confirm with the laboratory whether money is paid, the conversion rate, and which decisions count for payment. In hypothetical modules, describe amounts as outcomes in the stated situation.

When a comprehension check is included, an incorrect response receives an explanation beside that question, identifying the submitted answer and the correct answer. Other answers remain in the form. Ask the student to return to the example, locate the relevant rule, and redo the calculation before correcting the response. Avoid revealing answers to the entire class for an individual error, and do not interpret a comprehension mistake as an economic choice in the game.

la pública — Dos calendarios	Deuda pública,
la pública — Dos calendarios	Deuda pública, consumo y generaciones / Public Debt, Consumption and Generations — BEER Lab Seleccione una variante de este juego. Select a version of this game.
la pública — Cuatro perfiles: orden fijo	
la pública — Cuatro perfiles: orden aleatorio	
la pública — Referencia	
la pública — Sin préstamos privados	
la pública — Tasas de préstamo y depósito distintas	
la pública — Impuesto futuro incierto	
la pública — Contabilidad en valor presente	
la pública — Menor peso del consumo final	
la pública — Interés cero	
la pública — Dos generaciones: peso cero	
F la pública — Dos generaciones: peso intermedio	
la pública — Dos generaciones: peso completo	
la pública — Dos generaciones con reasignación	
d public debt — Two tax calendars	
d public debt — Four profiles: fixed order	
d public debt — Four profiles: random order	
d public debt — Baseline	
d public debt — No private borrowing	
d public debt — Different borrowing and deposit rates	
d public debt — Uncertain future tax	
d public debt — Present-value accounting	
d public debt — Lower final-consumption weight	
d public debt — Zero interest	
d public debt — Two generations: zero weight	
d public debt — Two generations: intermediate weight	
d public debt — Two generations: full weight	
d public debt — Two generations with rematching	

Figure 1.4: Demo contains only the configurations of this game family, in both languages and editions. Continuation 2/2.

Rooms

Current rooms:

Deuda pública — Básica
Public debt — Basic
Deuda pública — Extendida
Public debt — Extended

Figure 1.5: Rooms contains the four rooms for this family: basic and extended editions in Spanish and English.

Room: Public debt — Basic

Create a new session

Session Config

Consumption and public debt — Two tax calendars 

Number of participants

Individual decisions: any positive participant count is allowed.

Create

[Configure session](#) 

App sequence

public_debt_game (10 rounds)

Figure 1.6: Session-creation form for the basic English room. Check the selected configuration and participant requirement before creating the session.

Downloading a session with several rounds

For the instructor

Open **Data** within the relevant session and use that session's CSV or CSV-for-Excel download link. The general export has one row per participant and columns for the different rounds. Ten rounds do not necessarily mean ten rows per person. When the family provides its own long-format export, observations may instead be organized by participant and round or another unit defined by the design.

Save the original file together with the configuration name, session code, and active-round count. Check that columns for the last active round are present and that the data belong to the selected session. Empty cells can represent inactive rounds, a pending decision, or a field that does not apply to a role; do not automatically recode them as zeros. To compare learning, preserve round and group identifiers alongside participant identifiers. Share only anonymous aggregates in class, and record session incidents before constructing derived tables.

1.2 The question behind the activity

A cut in today's tax increases immediately available funds, but may also create a future obligation. This activity separates current liquidity, lifetime resources and chosen consumption. Participants encounter two calendars that finance identical public spending and choose how much to consume initially. The system records the remaining balance, applies interest and calculates final consumption. The comparison teaches a conditional proposition: under specified assumptions, replacing a current tax with public debt leaves a household's consumption opportunities unchanged.

The family is called Public Debt, Consumption and Generations. This descriptive name identifies the teaching product in the catalogue. Participant screens use action-based titles, such as Consumption and Saving Decisions. Introduce the theoretical hypothesis during discussion, after collecting choices. This arrangement allows students to use the budget information without being told the conclusion that their instructor intends to explain. The activity contains no demographic questionnaire and does not measure political support for a tax reform.

The baseline is individual. Each person represents a household that values consumption at two dates, without a human partner or a need to wait for other choices. Intergenerational extensions link two people: the first chooses consumption and a transfer, while the second receives resources and makes an unpaid transfer forecast. These institutions introduce nonnegative bequests and an announced weight on the recipient's consumption in the decision-maker's points. That weight is an explicit incentive, rather than an estimate of naturally occurring altruism.

1.3 Product facts

Purpose	Compare current and future tax finance and examine when saving or transfers offset a change in financing.
Experimental unit	An individual household in the baseline; a pair in intergenerational institutions.
Basic edition	Two calendars; four income profiles in fixed order; four profiles in random order.

Extended edition	Borrowing constraints, rate spreads, tax risk, accounting support, temporal preferences, zero interest and generation institutions.
Decision	Initial consumption; in pairs, an unpaid transfer forecast by the second generation.
Payoff	An affine logarithmic-utility score, summed across episodes; consumption is reported separately.
Participants	Any positive number individually; an even number in pairs, with at least four for rematching.
Suggested duration	45–75 minutes including instructions, choices and discussion.
Topics	Ricardian equivalence; public debt; consumption and saving; credit constraints; intergenerational transfers; tax uncertainty.
JEL	A22; C91; C92; D91; E21; H63.

1.4 Observable learning objectives

By the end, students should be able to reconstruct both budgets, discount a future obligation and explain negative saving. They should identify which assumption an extension changes and whether that change affects resources, preferences, information or the person receiving consumption. These tasks are more informative than a memorised definition of Ricardian equivalence. A satisfactory answer explains why constant public spending matters and why a borrowing restriction can change consumption even with identical present-value taxes.

For assessment, request three items: a budget identity, a numerical comparison and an interpretation of the report. The identity must include income, taxes, consumption and interest. In the comparison, students compute the saving difference while holding initial consumption constant. In the interpretation, they distinguish an observation from the class from a model prediction. They need no programming knowledge and should not be asked to explain administrative details of the application.

Adapt the depth to the course. In principles, use a timeline and the example with initial consumption of 80. Intermediate macroeconomics adds present value and the first-order condition. An experimental seminar can examine order, experience, observation dependence and the interpretation of incentives. All three routes use identical rules and accounts; the depth of explanation and the number of displayed statistics change. Avoid modifying the rules verbally to fit an explanation that assumes a different model.

For the instructor

Begin with two individual calendars and four test participants. Verify instructions, outcomes and completion before distributing class links. Reserve intergenerational pairs for a session with time to explain roles, transfers and waiting.

1.5 Preparation before students arrive

Check connectivity and enter the instructor interface. Select language and edition before creating links. A Spanish session does not become English because a browser uses English: language belongs to the configuration. Confirm participant count and mode. Individual configurations impose no divisibility requirement, whereas an incomplete pair cannot settle its episodes. Four participants are the minimum for rematching so that changing partners is possible. The default demonstration

size is a test arrangement, rather than the required class size.

Prepare a slide with the two dates and leave consumption blank. Do not display the optimum yet. Have a board or spreadsheet ready to add discounted income and subtract discounted taxes. Explain any real payment conversion that has been arranged. The application calculates performance ECU and does not itself establish a monetary conversion rate. If the activity is used solely for learning, announce that condition explicitly without promising payments that have not been organised.

In the test session, inspect one result with negative saving and another with a deposit. Verify that final consumption follows the announced equation and government borrowing does not appear as an unaccounted gift. Check that the report starts hidden and its reveal button waits for all active episodes. A completed first choice does not constitute a completed comparable block. Preserve other laboratory instances that are in use and create a new session when a fresh practice group is needed.

1.6 A minimum classroom sequence

Allow about ten minutes for context and rules, five for comprehension, ten for choices and twenty for discussion. Generations need another ten or fifteen minutes for roles, waits and additional blocks. Actual duration depends on class size and familiarity with budgets. The suggested duration is a lesson plan, rather than an automatic deadline that closes choices. Students can read the announced accounts before submitting their consumption.

Explain initially that relevant taxes and rates are announced and public spending is identical across baseline calendars. During choices, answer questions about rules rather than the consumption that maximises points. After completion, ask students to reconstruct a budget before revealing aggregate data. This step connects the report to known decisions. Starting with the class mean may lead students to confuse a behavioural regularity with an accounting identity. The discussion should distinguish what follows from algebra from what depends on optimisation.

Baseline budgets and predictions

2.1 Two dates and an identical public purchase

Let Y_1 and Y_2 denote initial and final income. Both equal 100 in the reference example. The government purchases goods worth $G = 20$ initially. The public service is identical in both calendars, so its constant benefit is omitted from the score. Under current tax finance, $(T_1, T_2) = (20, 0)$. Under public debt, $(T_1, T_2) = (0, 22)$ when $r = 0.10$. The 22 is principal plus interest, rather than an arbitrary extra charge.

The public creditor is external to participants. It lends 20 to the government, which purchases goods and repays 22 finally. A household may save or borrow with a separate external private bank at the same rate in the baseline. It is not required to buy government bonds. This separation avoids counting the government loan as household wealth while also allowing the household to spend it without a corresponding liability. The increase in household funds initially comes from the tax that is no longer collected at that date.

Initial consumption is C_1 , and the initial balance is

$$S = Y_1 - T_1 - C_1.$$

A positive balance is a deposit; a negative balance is a private loan. Final consumption is

$$C_2 = Y_2 - T_2 + (1 + r)S.$$

All final resources are consumed automatically. There is no third date or terminal saving value. Adding another saving choice at that point would introduce a decision without an economic destination in this model. Consequently, one economically meaningful choice determines both consumption amounts and the private financial contract connecting them.

2.2 An example without calculus

Suppose the household chooses initial consumption of 80. Under current tax it has $100 - 20 = 80$, saves zero and consumes 100 finally. Under public debt it has 100, saves 20 and obtains $100 - 22 + 1.10(20) = 100$ finally. The consumption pair 80 and 100 is feasible under both calendars. Saving differs by exactly 20. The household reserves the current tax cut for the future obligation, including the interest earned on its deposit.

Now choose $C_1 = 90$. Under current tax, saving is $80 - 90 = -10$ and final consumption is $100 - 11 = 89$. Under future tax, saving is $100 - 90 = 10$ and final consumption is $78 + 11 = 89$. Public and private debt are distinct: one calendar requires a private loan and the other a deposit. The same consumption pair remains feasible because both contracts use the same rate. This

example helps identify the mistaken belief that any negative saving means default or a violation of the rules.

Theoretical foundation

Calendar neutrality follows from identical feasible sets. Each admissible initial consumption under one calendar produces the same final consumption under the other. Identical optimal consumption then follows from this identity and unchanged preferences.

2.3 Present value and available wealth

Divide the final budget by $1 + r$ and rearrange:

$$C_1 + \frac{C_2}{1+r} = Y_1 - T_1 + \frac{Y_2 - T_2}{1+r} = W.$$

W measures net resources in initial-date units. It is not cash available today or a bank-account balance. Under current tax, $W = 80 + 100/1.10 \simeq 170.91$. Under public debt, $W = 100 + 78/1.10 \simeq 170.91$. The future obligation of 22 has present value 20. Discounting is necessary to compare quantities received or paid at different dates.

Both consumptions must be at least one to ensure positive consumption and a defined logarithm. With baseline parameters, maximum initial consumption compatible with $C_2 \geq 1$ is $W - 1/1.10 = 170$. The displayed interval is therefore 1 to 170. The interface accepts hundredths and rounds a computed upper bound down when necessary, ensuring decimal approximation cannot push final consumption below one. The lower bound is an announced technical restriction, rather than an estimated subsistence requirement.

An equal sum of incomes does not guarantee equal present value. Receiving more initially gives more resources in initial-date units than receiving that amount later when interest is positive. This distinction becomes visible in the basic income-profile battery. Neutrality is a comparison between tax calendars within a fixed profile. It does not imply that every profile, interest rate or preference weight produces the same optimal consumption. Keep these comparison dimensions separate when explaining the budget line.

2.4 Preferences and the benchmark solution

Individual utility is

$$J = \ln C_1 + \beta \ln C_2, \quad \beta = \frac{1}{1+r}.$$

The logarithm increases with consumption but has a declining slope: an additional unit matters more when consumption is low. The parameter β weights final utility and equals approximately 0.909091 in the baseline. It is an announced rule, not a preference inferred from participants' responses. Points equal $20 + 10J$, rounded to two decimals. This fixed positive affine transformation preserves the utility-maximising choice. Points do not add funds to the consumption budget.

Substitute $C_2 = (1+r)(W - C_1)$. The interior first-order condition is $1/C_1 = \beta/(W - C_1)$, giving

$$C_1^* = \frac{W}{1+\beta}, \quad C_2^* = \beta(1+r)C_1^*.$$

With the announced weight, both benchmark consumptions are approximately 89.52. Current-tax saving is $80 - 89.52 \simeq -9.52$; future-tax saving is $100 - 89.52 \simeq 10.48$. A saving difference of 20 accompanies a consumption difference of zero. The solution requires borrowing freedom up to the announced bound, certainty and identical private rates.

2.5 What the activity establishes

This identity does not imply that every deficit is neutral. Spending, incomes, rates and preferences are held constant, and the same household accounts for taxes at both dates. The model omits distortionary taxation, productivity changes, general-equilibrium effects and an endogenous interest rate. Changing spending removes the controlled comparison in the baseline. Assigning the future tax to another person without a compensating transfer also changes the decision problem.

Students may choose different consumption amounts despite identical feasible sets. That finding describes behaviour under this activity's information and incentives. It does not invalidate the accounting identity: choices may depart from the optimisation benchmark or reflect misunderstanding. Distinguishing a theoretical proposition from an empirical regularity is central to the discussion. It also prevents treating a numerical choice in a simplified environment as evidence of a participant's political preferences.

Checkpoint

Before revealing data, check that students calculate $22/1.10 = 20$, identify who repays the public loan and construct the same consumption pair using different balances. Return to the budget if one of these steps is missing.

Access, sessions and basic modes

3.1 From the portal to the correct room

The institutional starting point is <https://beer-lab.org/>. Open the teaching catalogue at <https://teaching.beer-lab.org/> and locate Public Debt, Consumption and Generations. Topics help find the activity through consumption, saving or public debt. JEL classifications supplement navigation: A22 covers undergraduate teaching; C91 and C92 distinguish individual and group experimental dimensions; D91 concerns behavioural decision making; E21 covers consumption, saving and wealth; H63 covers public debt. These classifications describe the academic content of the game.

The product domain is <https://publicdebt.beer-lab.org/>. Four rooms separate language and edition:

- Basic English:
https://publicdebt.beer-lab.org/room/public_debt_en
- Extended English:
https://publicdebt.beer-lab.org/room/public_debt_extended
- Basic Spanish:
https://publicdebt.beer-lab.org/room/deuda_publica_es
- Extended Spanish:
https://publicdebt.beer-lab.org/room/deuda_publica_extendido

Instructor access uses <https://publicdebt.beer-lab.org/demo>. A room provides session entry; it is not a public results page or a file containing personal responses.

Within this domain, Demo should display only this family and Rooms exactly four rooms, separated by language and edition. Each basic room offers three modes; each extended room offers a baseline plus ten extensions. No other family should appear. Check that a basic room starts with a basic configuration in the correct language. If unrelated games appear, stop preparing that session and contact laboratory coordination before distributing participant links. The visible room options provide a straightforward check without requiring technical knowledge.

3.2 Basic configurations and choice order

Two tax calendars contains two episodes. Each participant encounters current and future tax once, with income 100 at each date. Order is counterbalanced across people. Half begin with current tax and half with future tax, as far as the session size permits. With an odd number, order-group

sizes differ by at most one. Counterbalancing prevents one calendar from systematically occupying the position in which students have the least experience. It does not remove all learning from the second choice.

Four profiles: fixed order contains eight episodes arranged into four two-calendar blocks. Income profiles are (100, 100), (60, 140), (140, 60) and (80, 120). Incomes and rules remain constant across calendars within each block; their distribution across dates then changes between blocks. Students observe neutrality within a profile while present-value wealth changes across profiles. Total income equals 200 throughout, but its present value differs because earlier and later receipt are not equivalent when interest is positive.

Four profiles: random order uses the same profiles, with random block order and a randomly selected first calendar within each block. Everyone encounters each profile once with both calendars. Randomisation creates no new income and changes no public spending. Comparable scenarios remain adjacent within their block. This mode illustrates why a position-by-position chart can mix profiles and calendars and should not be interpreted as a common policy timeline for the whole class.

These modes are individual activities, rather than fixed-partner or rematching games. Repeating a comparison between dates does not create strategic interaction. Learning comes from different budgets and their consequences, without carrying saving between episodes. The choosing unit remains a person in the report. Several choices from the same person are not several independent experimental units. The number of recorded decisions can grow while the number of participants remains unchanged.

For classroom use, fixed order supports a shared explanation of income profiles. Random order helps discuss sequencing and encourages students to identify their current profile from the screen. Neither mode guarantees that a later mean measures improvement: profiles themselves change the consumption benchmark. Compare choices with the profile-specific benchmark or use matched calendar differences. This avoids mistaking an income-composition effect for a learning effect.

3.3 Create a session and distribute links

From Demo or the administrative room, select a configuration and enter participant count. Four is the suggested demonstration size. Check language in the descriptive label and confirm edition before creating links. The session panel then provides individual links and private monitoring and report access. Assign a distinct individual link to each student, or use the room's shared-entry mechanism if that is the planned classroom procedure. Do not project a complete list of personal links to the entire class.

A room may show an entry wait before the instructor activates a session. That entry screen differs from intergenerational settlement waiting. Once the individual activity begins, each student advances without synchronising choices with others. Avoid requiring everyone to press continue simultaneously: it introduces unnecessary synchronisation and can prevent careful budget reading. The instructor can monitor completion without displaying economic results that might affect choices still underway.

Individual links depend on the session database. Preserving it retains records; recreating it invalidates links tied to old sessions. Updating code does not require deleting the database to refresh screens. For a new practice group, create a new session and distribute its new links. Do not present previous responses as though they belong to the current class. If a student reopens the same link, check their existing progress before assigning another participant identity.

Create session

Tip: consider creating your session in a [Room](#), for more stable start links that are not opened accidentally.

Session Config

Consumption and public debt — Two tax calendars

Number of participants

4

Must be a multiple of 1

Create

[Configure session](#) 

App sequence

public_debt_game (10 rounds)

Figure 3.1: Select the English mode and confirm four test participants. The same procedure supports the actual class size.

oTree Demo Sessions Rooms Data Server Check Logout

Consumption and public debt — Two tax calendars: session

 Edit  Links  Monitor  Data  Payments  Report

 Description

You can either use the session-wide link, persistent links, or single-use links.

Session-wide link

If it's not practical for you to use a Room, you can distribute the following link to participants.

[Shared session link](#)

Persistent links

For a real session, you should use a [room](#).

Single-use links

Below are single-use links.

P1	Participant link
P2	Participant link
P3	Participant link
P4	Participant link

Figure 3.2: The session panel provides participant entry and private session navigation. Test links use impersonal labels in this capture.


3.4 Local testing and product routes

The package includes a local instance at <http://localhost:8090/demo>. Localhost is available on the computer running the service. It is not an Internet address for students using other devices. Classroom access from students' own computers or phones requires the deployed product domain. The canonical routes describe the intended product structure; local captures verify that structure in the test instance. A PDF or a Git commit does not establish that a public domain is active.

Participant screens and actions

4.1 Welcome and comprehension

The welcome explains announced income and obligations, fresh episode resources and experimental points. It presents general laboratory rules and requests acceptance before proceeding. Headings describe choices and outcomes, without internal production codes or the hypothesis discussed later. Any real incentive arrangement must be explained consistently with the written budgets. An oral explanation should clarify the interface, rather than substitute a different payment or borrowing rule.



  Episode 1 / 2

Consumption and Saving Decisions

You will make decisions with announced income and obligations at two dates. Every episode starts with fresh resources. Your choices generate experimental points (ECU).

Laboratory rules

- Decide individually and keep your personal link private.
- Read the rules; there is no deception and no demographic questionnaire.
- Results will be discussed in aggregate after completion.

☐ I have read and accept the rules.

Continue

Figure 4.1: The welcome obtains acceptance and explains the task. Participants keep their own session links private.

Comprehension has two questions. One checks that the government loan is repaid through future taxes and an external fund when announced. The other checks that the initial balance subtracts taxes and consumption from income. Incorrect responses retain the participant on the page and



Episode 1 / 2

Instructions

Resources and unit of account

Pre-tax income in the displayed scenario: period 1, 100 ECU; period 2, 100 ECU. Available after the period-1 tax: 80 ECU.

ECU stands for Experimental Currency Units. It is the unit of account for this activity.

This session does not convert ECU into real money.

The government buys 20 ECU of goods. With a current tax it collects 20 now and 0 later. With a future tax it borrows 20 externally now and collects principal plus interest later. The public service is identical and does not change your points. Private loans settle with an external bank; you are not required to buy government bonds.

Choose initial consumption. The balance equals initial income minus initial tax minus consumption. Final consumption equals final income minus final tax plus the balance with interest. All final resources are consumed; there is no value for resources after the episode.

Each person decides for a household. There are no pairs and no waiting for other decisions. You may save or borrow. Interest applies to the balance; final consumption must be at least 1.

In the individual activity, points = $20 + 10 \times [\ln(\text{initial consumption}) + \beta \times \ln(\text{final consumption})]$. $\ln$ is the natural logarithm; β is displayed before you choose. Points across episodes are added. Consumption and saving use budget ECU; points are a performance index and do not add funds to the budget.

Baseline

Episode result

Individual example: with income 100 and 100, current tax 20, a 10% rate and initial consumption 80, saving is 0 and final consumption is 100. With a future tax of 22 and initial consumption 80, saving is 20 and final consumption is also 100.

Continue

Powered by [oTree](#)

Figure 4.3: Current instructions for the basic edition: initial resources, decision rules, and experimental units. Continuation 2/2.

request rule review. Neither question asks for point-maximising consumption. Their purpose is to establish understanding of obligations and resources, rather than test whether the student remembers a supplied solution.



Review the marked questions. Each shows your answer and an explanation; your other answers are retained.



Episode 1 / 2

Check the rules

What happens to the government loan?

It disappears without repayment.

Your answer: It disappears without repayment.. The correct answer is It is repaid using future taxes and, if announced, an external fiscal fund.. Review loan repayment and the calculation of the balance.

How is the initial balance calculated?

Initial income minus initial tax minus initial consumption.

Continue

Powered by [oTree](#)

Figure 4.4: Current comprehension check with an incorrect answer, its specific explanation, and retained responses.

4.2 The individual decision screen

The screen shows the tax calendar, both incomes, taxes, rates and final-consumption weight. The consumption field states its allowed interval. A participant enters initial consumption to hundredths and continues. Values outside the interval are rejected. The upper bound preserves positive final resources and, when borrowing is prohibited, also prevents spending beyond current funds. Students can identify the binding constraint from the announced quantities without a label suggesting the desired answer.

Saving is not a separate independent choice. Once consumption is fixed, $S = Y_1 - T_1 - C_1$ determines the balance. Separate unrestricted fields for saving and consumption could permit incompatible choices and require a distracting reconciliation step. Final consumption also needs no separate choice because there is no terminal saving destination. The design leaves one meaningful economic allocation and displays its consequences. Forecasts occur only in generations, where the recipient understands the environment but cannot control the transfer.




Episode 1 / 2

Choose your initial consumption

Current tax

Initial income	100.0	Final income	100.0
Initial tax	20	Expected final tax	0

Deposit / borrowing rate: 10% / 10%

Final-consumption weight β : 0.909091

You may save or borrow. Interest applies to the balance; final consumption must be at least 1.

Allowed interval for initial consumption: 1–170.0 ECU

Initial consumption (ECU)

Initial resources after tax: 80.0 ECU

Continue

Figure 4.5: Read income, taxes and interest before entering consumption. The announced interval is part of the budget rules.

Present-value accounting adds a box containing consolidated resources and a discounting explanation. It makes an operation available from the baseline information easier to perform, without changing income, taxes or rates. A comparison can be useful for teaching, but showing the support screen can itself generate learning. In a sequential classroom demonstration, that learning is intended. It also means the comparison should not automatically be described as a research estimate of a pure information effect.

4.3 Outcomes and transition

Outcomes separate initial consumption, final consumption, balance and points. They also display realised final tax, government borrowing, repayment and the fiscal-fund contribution. A negative balance is a private obligation to the external bank. A negative fund contribution means the fund receives surplus revenue, rather than an additional score penalty for the household. Points depend on consumption utility; taxes and interest have already entered consumption amounts. Deducting them again from points would count the same effect twice.

Current tax

Initial consumption (ECU)	Final consumption (ECU)
89.52	89.53
Initial saving or transfer (ECU)	Episode points (ECU)
-9.52	105.80
Realised final tax	0.00
External government borrowing	0.00
Government repayment	0.00
External fiscal-fund contribution (negative: receives surplus)	0.00

Figure 4.6: Budget quantities and the performance score appear separately. Reconstruct final consumption from the displayed balance.

Every episode ends without capital or private debt carrying into the next. Accumulated points record performance and cannot be reused as a consumption endowment. The next screen may show a different calendar or profile. Explain this transition to students familiar with a growth game: here saving connects two dates within one episode, whereas growth investment changes a productive state that persists. Confusing these transitions changes both choice interpretation and the meaning of a position-by-position chart.

4.4 Roles and intergenerational waiting

Pairs receive a role page at each episode's start. The first generation chooses consumption and leaves a nonnegative balance. The second sees the same budget rules and forecasts the initial transfer. Its prediction changes no consumption, earns no reward and is not shown to the chooser before the choice. After both responses, a wait page synchronises the pair and the system settles consumption and points. No participant negotiates, rejects or returns a transfer in this institution.




Episode 1 / 10

Your role

First generation

The pair remains together; roles switch after each two-episode block.

The first generation chooses own consumption and transfers the balance to the second. First-generation points = $20 + 10 \times [\ln(\text{own consumption}) + \alpha \times \beta \times \ln(\text{second-generation consumption})]$. Second-generation points = $20 + 10 \times \ln(\text{received consumption})$. The weight α is an announced scoring rule. The second generation makes an unpaid forecast and cannot reject or return the transfer.

Current tax

Initial income	100.0	Final income	50
Initial tax	20	Expected final tax	0

Deposit / borrowing rate: 10% / 10%

Final-consumption weight β : 0.909091 · Announced second-generation weight α : 1

You cannot borrow. The unconsumed balance is nonnegative saving or a transfer.

Continue

Figure 4.7: The role screen identifies the transfer decision-maker and the announced recipient weight. Check who pays the final tax.

Partners and roles remain unchanged across both calendars of a comparable block. After a block, fixed pairs exchange roles. Under rematching, new partners are drawn and positions determine the role in the next block. Random assignment does not guarantee exactly five episodes in each




Episode 1 / 10

Anticipate the transfer

Future tax

Initial income	100.0	Final income	50
Initial tax	0	Expected final tax	22.0

Deposit / borrowing rate: 10% / 10%

Final-consumption weight β : 0.909091 · Announced second-generation weight α : 1

You cannot borrow. The unconsumed balance is nonnegative saving or a transfer.

Predict the balance that the first generation will transfer, before you observe it. Your forecast does not change points or the budget.

Expected initial transfer (ECU)

Continue

Figure 4.8: The second generation forecasts an initial transfer without a reward. The response describes an expectation and does not authorise a different allocation.




Episode 1 / 10

Wait for the other person

The result is calculated when the first person has chosen and the second has recorded the forecast. Keep this page open.

Figure 4.9: Waiting synchronises only the pair. Delay indicates a missing response and reveals no economic choice.

role for every individual. Accordingly, comparisons use the same chooser within a block, without assuming identical lifetime role exposure in a rematched session.

What participants will see

Participants receive information needed to act: calendar, income, rates, limits, role and scoring rule. The recipient makes a forecast and observes an outcome. There is no acceptance, negotiation or reverse-transfer decision.

4.5 Explicit completion

The final page confirms all choices are complete, displays accumulated points and shows a personal history. There is no continue button inviting students into inactive episodes. Keep this page available for discussing personal budgets without projecting identifiers. The instructor reveals aggregate outcomes through the private report. Early finishers can reconstruct a block on paper while waiting for discussion; they need not submit further answers or search for another completion page.



Figure 4.10: Completion lists both calendars and accumulated points. Compare consumption and balances within a block before comparing different modes.

Extensions: credit, risk and information

5.1 Compare one mechanism at a time

The extended edition includes an individual baseline and ten extensions. Six retain a household as the choosing unit: no private borrowing, different borrowing and deposit rates, uncertain future tax, present-value accounting, lower final-consumption weight and zero interest. Four use pairs: zero, intermediate and full recipient weights, and full weight with rematching. Their selection reflects useful changes in constraints, preferences or information, rather than a predetermined numerical limit on configurations.

Do not run every extension in a single lesson. Select a question and compare conditions that make it informative. For liquidity, compare the baseline and no-borrowing condition while keeping income 100 and 100. For intergenerational transmission, compare announced weights with incomes 100 and 50. For information, preserve economic parameters and change accounting support. Comparing activities with different experimental units can be conceptually useful but does not estimate a single isolated effect.

Before showing a comparison, ask which quantity or rule changes. An interest-rate change affects the budget slope; a recipient weight affects the objective; accounting support affects presentation; nonnegative saving restricts feasible choices. Students often describe all such changes simply as a different treatment. Naming the altered mechanism gives them a concrete way to derive predictions and prevents attributing every difference to the tax calendar itself.

5.2 No private borrowing

Prohibiting loans requires $S \geq 0$, giving $C_1 \leq Y_1 - T_1$. Under current tax the maximum is 80; under future tax it is 100. The unconstrained benchmark is approximately 89.52 in both calendars. A constrained household cannot reach it when paying 20 today but can when that tax is deferred. Optimal current-tax initial consumption becomes 80, while future-tax consumption remains approximately 89.52. Equal present-value wealth is insufficient for identical consumption opportunities when future funds cannot be brought forward through borrowing.

The benchmark consumption increase is approximately 9.52. Saving moves from zero to 10.48, increasing by less than the tax cut of 20. The compensation proportion is approximately $10.48/20 = 0.524$. These quantities are parameter-based predictions, not an estimated marginal propensity to consume or a spending multiplier. Public spending is constant, so the mechanism operates through liquidity and household feasibility. The present-value identity continues to hold even though a

segment of the budget line becomes inaccessible.




Episode 1 / 2

Choose your initial consumption

Current tax

Initial income	100.0	Final income	100.0
Initial tax	20	Expected final tax	0

Deposit / borrowing rate: 10% / 10%

Final-consumption weight β : 0.909091

You cannot borrow. The unconsumed balance is nonnegative saving or a transfer.

Allowed interval for initial consumption: 1–80.0 ECU

Initial consumption (ECU)

Initial resources after tax: 80.0 ECU

Continue

Figure 5.1: Consumption is limited by current resources after tax. Compare the upper bound with the unconstrained benchmark of 89.52 before interpreting choices.

Ask students to identify the excluded consumption point. If someone says wealth rose, recalculate W under both calendars. The restriction removes choices previously financed through loans without changing present-value taxes. In the report, inspect condition-specific benchmark distance and upper-bound choices. A boundary choice is consistent with a binding restriction, but does not alone demonstrate that its chooser solved the optimisation problem. Correct inference requires the announced constraint and objective, not the boundary label in isolation.

5.3 Different borrowing and deposit rates

This extension keeps a 10% deposit rate and charges 25% on private borrowing. Government debt is still repaid at 10%. For $S \geq 0$, final consumption adds $1.10S$; for $S < 0$, it adds $1.25S$. The budget changes slope at $C_1 = Y_1 - T_1$. A single common discount rate no longer consolidates every

choice into an equivalent straight line. Present value at the deposit rate remains an accounting reference but cannot alone represent opportunities involving more expensive loans.

Under current tax, equivalent resources in the borrowing region are $80 + 100/1.25 = 160$. Dividing by $1 + \beta$ gives approximately 83.81, above 80 and therefore within that region. Under future tax the optimum of approximately 89.52 lies below current available funds of 100 and uses a deposit. Additional saving no longer fully offsets tax timing because avoiding expensive private credit changes the relative price of consumption. The benchmark calculation also evaluates the slope-change point and feasible endpoints.




Episode 1 / 2

Choose your initial consumption

Current tax

Initial income	100.0	Final income	100.0
Initial tax	20	Expected final tax	0

Deposit / borrowing rate: 10% / 25%

Final-consumption weight β : 0.909091

You may save or borrow. Interest applies to the balance; final consumption must be at least 1.

Allowed interval for initial consumption: 1–159.2 ECU

Initial consumption (ECU)

Initial resources after tax: 80.0 ECU

Continue

Figure 5.2: Both rates are announced before choosing. A negative balance uses the borrowing rate and a positive balance uses the deposit rate.

Use the contrast to explain why discounting everything at 10% can be insufficient. This institution introduces neither default risk nor a change in the government's borrowing rate. Rates are exogenous and known. Class outcomes cannot establish that banks increased rates in response to public debt. The report separately displays net interest and a benchmark that accounts for the appropriate rate on each side of the balance boundary. A difference between calendars can therefore be rational under the altered opportunity set.

5.4 Uncertain future tax

The current-tax calendar remains 20 now and zero later. Under future tax, the final obligation is 11 or 33 with equal probability. Its mean of 22 matches the certain baseline obligation. The government always repays 22. An external fiscal fund contributes 11 when revenue is 11 and receives an 11 surplus when revenue is 33. Household tax plus the signed fund contribution therefore equals repayment in every state. The fund does not directly change a household's score.

This institution changes the final-tax distribution as well as its timing. Expected wealth at the 10% rate is equal across calendars, but contingent budgets differ. Describing it as a neutrality test under identical risk would be incorrect. Consumption is chosen before the realised tax is disclosed. The allowed interval ensures $C_2 \geq 1$ even with tax 33, making maximum initial consumption 160 in the risky calendar rather than 170 in the certain baseline. This bound protects logarithmic utility in both announced states.

The benchmark maximises expected utility,

$$\ln C_1 + \frac{\beta}{2} [\ln C_2^{11} + \ln C_2^{33}].$$

It does not substitute the logarithm of expected consumption. Because the logarithm is concave, averaging utilities differs from evaluating utility at an average quantity. Realised points depend on the observed state, whereas benchmark loss uses expected utility before the draw. A fortunate tax outcome cannot retrospectively turn any chosen consumption into the best ex ante allocation.

For an example, hold initial consumption at 80. Future consumption is 111 under tax 11 and 89 under tax 33, with mean 100. The current-tax calendar delivers a certain 100. Their expected consumption matches, but expected logarithmic utility is lower for the spread around 100. This comparison separates an accounting equality from risk valuation. The optimal response need not be inferred from a single realised score, since the same decision can receive different realised points across the two states.

Attention

The risky extension compares a certain calendar with another containing uncertainty and a contingent fiscal fund. Expected tax remains constant, but risk changes. An observed difference cannot by itself identify the pure effect of postponing taxes.

5.5 Present-value accounting support

This extension adds consolidated wealth and a discounting explanation. It preserves baseline income, spending, rates, temporal weight and credit access. The economic information is already present in the baseline, but its organisation changes. The optimisation prediction remains equal consumption with saving differing by 20. Better performance is compatible with reduced calculation difficulty or greater attention to future liabilities. Additional measurements would be needed to distinguish these explanations.

Use this contrast to show that computational mistakes can resemble preference differences. Ask which operation the box facilitates and whether the student could perform it from the baseline information. Examine benchmark distance and utility loss rather than whether saving simply increased. Higher saving is not always closer to the optimum: someone already saving too much can improve by consuming more initially. Keeping the performance criterion explicit avoids equating any reduction in consumption with better reasoning.




Episode 1 / 2

Choose your initial consumption

Future tax

Initial income	100.0	Final income	100.0
Initial tax	0	Expected final tax	22.0

Deposit / borrowing rate: 10% / 10%

Final-consumption weight β : 0.909091

You may save or borrow. Interest applies to the balance; final consumption must be at least 1.

With a future tax, the amount will be 11 or 33 ECU, each with probability 50%. An external fiscal fund covers the difference from the loan repayment of 22 (or receives the surplus). Choose before the draw is revealed. The allowed interval ensures positive final consumption in both states.

Allowed interval for initial consumption: 1–160.0 ECU

Initial consumption (ECU)

Initial resources after tax: 100.0 ECU

Continue

Figure 5.3: The future-tax screen states both obligations, probabilities and the fiscal fund's role. The draw is disclosed after the choice.



Episode 1 / 2



Episode 1 / 2

Choose your initial consumption

Current tax

Initial income	100.0	Final income	100.0
Initial tax	20	Expected final tax	0

Deposit / borrowing rate: 10% / 10%

Final-consumption weight β : 0.909091

You may save or borrow. Interest applies to the balance; final consumption must be at least 1.

Present-value accounting

Discount final resources at the deposit rate. In the equal-rate baseline, adding initial net resources and discounted final resources summarises the two-date budget. With different rates this is an accounting reference, rather than a single budget line.

Present value of expected resources (ECU): 170.91

Allowed interval for initial consumption: 1–170.0 ECU

Initial consumption (ECU)

Initial resources after tax: 80.0 ECU

Continue

Figure 5.5: The extra box displays resources in present-value terms. It organises existing information without adding income or changing taxes. Continuation 2/2.

5.6 Temporal weight and zero interest

Lower final-consumption weight sets $\beta = 0.5$, with 10% symmetric private rates. The initial benchmark is $W/1.5 \simeq 113.94$. The household accepts lower final consumption because the scoring rule gives that date less weight. Feasible sets still match across calendars, and optimal saving still differs by 20. This comparison separates a preference for earlier consumption from neglect of future tax obligations. Impatience and tax-timing neutrality can coexist within the model.

Zero interest sets $r = 0$ and $\beta = 1$. Future tax repays public principal of 20, rather than 22. Present-value resources equal 180 and both benchmark consumptions are 90. Neutrality no longer requires numerical discounting but still requires accounting for a future obligation. Use this case first when students have not learned present value, then restore the 10% rate. Its instructional simplicity does not make zero interest a general theory of monetary economies.

Transfers between generations and research evidence

6.1 One pair and two consumption recipients

The first generation receives 100 initially and the second receives 50 finally. The government finances the same purchase of 20 with a current tax or external debt repaid through a future tax of 22. The first chooses C_1 and deposits a nonnegative transfer $B = Y_1 - T_1 - C_1$. The second consumes $C_2 = Y_2 - T_2 + 1.10B$. Reverse transfers are unavailable. The external private bank pays deposit interest, while the public creditor receives repayment from government. These contracts are separately accounted for.

First-generation points are $20 + 10[\ln C_1 + \alpha\beta \ln C_2]$; second-generation points are $20 + 10 \ln C_2$. The announced α is 0, 0.5 or 1. The second generation forecasts B without reward; aggregate absolute forecast error is reported. This structure turns fiscal transmission into a distributive decision with intertemporal budgets. It does not replicate a stationary overlapping-generations economy or an infinite bequest chain. Each two-calendar block begins again with fresh resources.

6.2 Full weight and an operative bequest

With $\alpha = 1$, the first generation values both utilities using the baseline temporal weight. Consolidated resources are $W = 80 + 50/1.10 \simeq 125.45$. Benchmark initial consumption is $W/(1 + \beta) \simeq 65.71$, and final consumption also equals approximately 65.71. Under current tax, the transfer is $80 - 65.71 \simeq 14.29$. Under future tax, it is $100 - 65.71 \simeq 34.29$. Their difference of 20 accumulates to 22 and offsets the future tax. Because the first transfer is already positive, the nonnegative-bequest restriction does not prevent adjustment here.

An operative bequest means resources would be transferred even before the financing change. Merely allowing a transfer is insufficient: if its optimum is zero, the lower bound may prevent adjustment in one direction. Ask students to calculate both transfers and recipient consumption. Equal recipient consumption does not imply equal nominal bequests. Display all three quantities to prevent the mistaken conclusion that public debt prompted no response at all.

6.3 Zero and intermediate weights

With $\alpha = 0$, the first generation maximises own consumption in its scoring rule. It consumes 80 under current tax and 100 when tax is deferred. Benchmark transfers are zero in both cases. The second receives 50 or 28. Financing shifts consumption across people because the recipient




Episode 1 / 10

Choose your initial consumption

Current tax

Initial income	100.0	Final income	50
Initial tax	20	Expected final tax	0

Deposit / borrowing rate: 10% / 10%

Final-consumption weight β : 0.909091 · Announced second-generation weight α : 1

You cannot borrow. The unconsumed balance is nonnegative saving or a transfer.

Allowed interval for initial consumption: 1–80.0 ECU

Initial consumption (ECU)

Initial resources after tax: 80.0 ECU

Continue

Figure 6.1: The first generation chooses consumption from current funds and leaves a nonnegative transfer. The announced weight identifies the scoring rule.

of the tax cut has no point-based reason to compensate the future taxpayer. The family is no longer a dynasty fully internalising both consumptions. This change in recipients and preferences is economically substantive.

With $\alpha = 0.5$, unconstrained first-generation consumption would be 86.25. Under current tax it cannot exceed 80 because a negative bequest is prohibited. It chooses 80 and transfers zero. Under future tax it chooses 86.25 and transfers 13.75; recipient consumption becomes $28 + 1.10(13.75) = 43.125$. Compensation is partial, with ratio $13.75/20 = 0.6875$. This institution illustrates a transfer boundary rather than a linear interpolation between zero- and full-weight outcomes. The constraint can bind under one calendar but not the other.

A positive transfer under zero weight may be compatible with consideration for the recipient, norms or mistakes. It does not distinguish those explanations. Nor does a transfer increase under full weight estimate natural altruism, since recipient consumption is explicitly rewarded. Separate the announced scoring criterion, any additional private preferences and the recorded decision. Reported benchmarks describe the incentivised objective and are not structural estimates of private social preferences.

6.4 Fixed pairs and between-block rematching

The three weight institutions keep pairs for ten episodes, organised as five two-calendar blocks. Roles alternate after each block. Full weight with rematching preserves scoring but draws new partners at each block's start. Both calendars within a block keep the same people and roles. A within-block consumption change therefore describes the same chooser's response, rather than confounding a tax change with a new decision-maker or recipient halfway through the comparison.

Rematching contrasts repeated exposure to one recipient with exposure to changing recipients. The full-weight material benchmark remains complete compensation in either matching regime. Observed differences may relate to learning, expectations or social considerations, but classroom outcomes do not automatically identify each mechanism. Experience travels with a person who changes partners. Explain what remains fixed, what is redrawn and which elements of personal history are not erased by rematching.

6.5 Relationship to research

[Barro \(1974\)](#) develops the theoretical neutrality argument with relevant intertemporal links. This product adopts a transparent two-date budget and external contracts, rather than reproducing its general-equilibrium environment. [Cadsby and Frank \(1991\)](#) study overlapping generations experimentally; their published abstract highlights approximate predicted outcomes with positive bequests and the role of nonnegative-transfer restrictions. That evidence motivates contrasting operative bequests with transfer boundaries, without determining this guide's calibration.

[Meissner and Rostam-Afschar \(2017\)](#) examine learning in a lifecycle consumption experiment. Their abstract reports average departures from equivalence and changes in behavioural rules with experience. This product shortens the horizon to make accounting visible. [Ricciuti and Di Laurea \(2003\)](#) examine liquidity constraints and income uncertainty in generations. Here uncertainty concerns taxes and uses a reconciled external fiscal fund, a distinct teaching adaptation. Outcomes from those studies are not presented as observations from any future classroom session.

Reading the report and leading discussion

7.1 Private access and controlled revelation

Open the report from the session panel using instructor access. It is neither a public room nor a participant's personal link. Initially it displays progress, pending records and completed decision count, while economic outcomes remain hidden. Revelation is enabled only after every active episode is complete, including both calendars in each block. Reloading hides results again, leaving the instructor to choose when to reveal. This allows preparation without projecting outcomes while classmates are still choosing.

Aggregates explicitly traverse the active episodes for the configuration. Entering the report through the last available technical period must neither exclude the first episode nor add inactive ones. With four participants in Two tax calendars, the final report contains eight complete decisions and four matched blocks. A ten-episode generation extension with four participants contains forty participant records, divided into twenty consumption choices and twenty recipient records. These counts answer different questions and should not be substituted for one another.

Global revelation prevents early finishers' aggregate choices from being projected to those still deciding. It does not remove personal feedback after an episode, which is part of announced learning. Nor can the application prevent conversations outside the interface. Organise quiet decision making if that is intended. None of these extensions adds unrestricted communication, and transfer forecasts are not disclosed to the chooser before their consumption decision.

7.2 Calendar comparison and complete blocks

Calendar counts include completed consumption choices. In pairs, each first-generation choice counts once per pair and episode. Calendar means describe complete records in that condition. Initial-consumption change instead uses only blocks containing both calendars for the same chooser. It calculates $C_1^{future} - C_1^{current}$ within each block, then averages these differences. It does not subtract two means based on different incomplete samples.

Saving or transfer change follows the same procedure: $S^{future} - S^{current}$. With identical initial income and spending of 20 within blocks, accounting gives $\Delta S = 20 - \Delta C_1$. The compensation proportion divides mean balance change by 20. One means an increase of 20 in the balance, while zero means no increase. Values can be negative or exceed one if choices move accordingly. The index is a magnitude, rather than a frequency or probability, so its range is not restricted to the unit interval.




Figure 7.1: Before completion, progress is visible and revelation is disabled. One finished decision is not a complete comparable block.




Figure 7.2: Completion enables the button but keeps outcomes hidden until the instructor reveals them.

Tax-calendar comparison

Current-tax decisions	4
Future-tax decisions	4
Initial consumption (ECU) — Current tax	77,615
Initial consumption (ECU) — Future tax	77,615
Complete comparable blocks	4
Initial-consumption change: future – current	0,0
Saving or transfer change: future – current	20,0
Balance change / 20 (compensation reference)	1,0
Mean distance from benchmark consumption	11,909
Mean expected-utility loss	0,038
Choices at the upper bound	0

Only completed active episodes enter aggregates. In pairs, each first-generation choice counts once; second-generation points are reported separately. Changes use only blocks where the same decision-maker completed both calendars. The compensation index is descriptive: a reference of 1 requires comparable calendars; uncertainty and different rates change the economic problem. The curve mixes calendars when order is counterbalanced and does not identify a causal effect.

Figure 7.3: Read calendar counts, matched-block count and within-block differences together. The denominator determines the comparison represented by a mean change.

Consider four illustrative blocks, unrelated to observed classroom data. Current-tax consumptions are 80, 90, 100 and 85; future-tax consumptions are 80, 95, 100 and 90. Differences are 0, 5, 0 and 5, averaging 2.5. Mean balance change is therefore 17.5 and compensation is 0.875. This does not mean 87.5% of students are Ricardian. It summarises an average saving response. Return to choices, blocks and the histogram to discuss heterogeneity.

The distinction matters especially in partial sessions. Suppose most current-tax episodes are complete but fewer future-tax episodes are complete. Calendar means may then describe different subsets, whereas a matched difference uses only people with both choices. Completion gating postpones classroom revelation, and matched denominators make the stored comparison defensible when inspected administratively. Neither feature converts a small convenience class into a representative sample of households.

7.3 Benchmarks, distances and utility

Mean benchmark distance averages $|C_1 - C_1^*|$ over complete choices. Each benchmark incorporates that condition's parameters, constraints and announced preferences. Under no borrowing and current tax, the reference is 80, rather than 89.52 as though loans were allowed. In intermediate-weight generations, a transfer boundary can determine the optimum. In the risky calendar, the benchmark maximises expected utility before the state is known, so a favourable tax draw does not retrospectively define the best initial choice.

Mean expected-utility loss subtracts each choice's expected utility from the best feasible expected utility. Its unit is logarithmic utility, rather than consumption ECU or government funds. Do not deduct it from consumption or a fiscal account. Equal consumption errors can generate different utility losses because marginal utility and curvature differ. Distance is easy to communicate, while utility loss helps distinguish a visibly large deviation from a costly one within the announced criterion.

Upper-bound counts use a tolerance of approximately one hundredth. Under perfect credit, consuming nearly the maximum leaves low final consumption and is generally unattractive with utility at both dates. Under zero recipient weight without borrowing, the upper bound is the first generation's optimum. The same boundary label therefore has different meanings across conditions. Relate it to the specific objective before describing participants as impulsive or rational. A boundary observation alone is not a behavioural diagnosis.

7.4 Distribution and evolution

The histogram counts initial-consumption decisions in 40-ECU intervals. The first includes minimum consumption of one. Endpoints are assigned once: a choice of 80 enters the interval ending at 80. Bar heights represent frequencies, not compensation percentages. Their counts sum to complete consumption choices. Recipient records are not added as if recipients had chosen initial consumption, which would duplicate observations from a single allocation.

Evolution shows mean initial consumption and a table containing position, decision counts, calendar composition, final consumption and balance. The horizontal axis is episode position, rather than years in an actual economy. Counterbalanced positions mix calendars, and the random battery can also mix income profiles. The line helps identify missing records or position-related patterns, but does not identify a causal fiscal-policy trajectory. Use complete-block differences for calendar comparisons.

Initial-consumption distribution




Figure 7.4: The histogram displays concentration and heterogeneity. Counts represent decisions, rather than unique participants in a multi-episode session.

Evolution by episode




Figure 7.5: The table accompanies the curve with counts and episode means. Inspect calendar composition before interpreting any position.

7.5 External accounts and recipient outcomes

The finance section sums public principal, public repayment, signed fiscal-fund contributions and net private interest. Principal and repayment reconcile with spending and taxes in each calendar. Under risk, realised tax plus fund contribution equals repayment. Private interest can be negative when households pay loan interest. A net zero does not mean nobody saved or borrowed. These totals reconstruct external flows without modelling a commercial bank's reserve or regulatory-capital constraints.

Generations also report mean second-generation points and recipient-record count. Mean absolute transfer-forecast error compares the unpaid forecast with the realised initial balance, in initial budget ECU. It is neither a direct trust measure nor a precision reward. The chooser sees no forecast before deciding. Use the statistic to ask how expectations formed under different weights, while recognising that the application does not elicit a complete probability distribution of beliefs.

Cumulative external accounts

External government borrowing	200,0
Government repayment	220,0
External fiscal-fund contribution (negative: receives surplus)	0,0
Net interest received by households	48,58
Mean second-generation points (n=20)	61,85
Mean absolute transfer-forecast error	14,29

Figure 7.6: Public accounts are separated from recipient points and forecast errors. Recipient outcomes state their corresponding record count.

Feedback script

Start with three questions: how many complete blocks are compared, what balance difference was predicted in this condition, and which quantity changed in our data? Then request a budget explanation and a compatible behavioural explanation. Keep them separate until evidence can distinguish them.

Lesson plans, exercises and data use

8.1 A 45-minute introductory lesson

Use the first eight minutes for the two dates, spending of 20 and repayment of 22. Present initial consumption of 80 and let students calculate both balances. Allocate five minutes to rules and comprehension, seven to individual choices, fifteen to the report and ten to a transfer exercise. Early finishers should reconstruct their budgets without discussing decisions with others. Preserve enough discussion time to move from numerical results to the assumptions that explain them.

Begin discussion with balances, then consumption. Ask someone to explain how a current-tax household can borrow without defaulting. Ask what would happen with different public and private rates, allowing a qualitative answer before calculation. Close by listing baseline assumptions and introducing a case that violates one. The objective is to know when the proposition applies and which accounts must be checked, rather than attach a permanent behavioural label to each student.

If the class has little experience with discounting, start with zero interest. Compute a future obligation of 20 and consumption 90 at both dates before moving to the 10% reference. This sequence isolates the accounting idea from the arithmetic of present value. Explicitly announce the new rate when changing conditions: otherwise students may interpret the future tax of 22 as an unrelated charge. A controlled comparison should leave its changed parameter visible.

8.2 A 75-minute intermediate lesson

Reserve fifteen minutes for the consolidated budget and first-order condition, ten for baseline choices, fifteen for a credit or information extension, twenty for report comparison and fifteen for exercises and evidence. If selecting generations, replace the individual extension with roles and paired blocks. Explain why the recipient has no reverse-transfer choice. Do not add informal negotiation that changes the announced institution or ask recipients to compensate a chooser outside the application.

The income-profile battery demonstrates that moving income across dates can change W while total income stays constant. Compare $(60, 140)$ and $(140, 60)$, then check neutrality within each profile. This separates a property across tax calendars from a property across income profiles. Changing both at once can produce a consumption difference driven by wealth rather than financing. A table with one block per profile keeps the comparison answering each question visible.

For a longer seminar, assign teams a different report statistic. One team reconstructs compensation, another explains benchmark distance, another reconciles external accounts and another evaluates dependence across observations. Each team should state numerator, denominator and units before offering an interpretation. This arrangement encourages scrutiny of the report without requiring

every student to learn administrative export details. It also shows that multiple statistics can describe the same choices without measuring the same mechanism.

8.3 Exercises with answer guidance

8.3.1 An apparently larger obligation

A student claims that collecting 22 later must be more costly than collecting 20 now. Evaluate the claim at a 10% interest rate. The answer compares $22/1.10$ with 20 and finds equality. Then prohibit private borrowing. Present-value taxes remain equal, but feasible consumption can differ. A complete answer distinguishes tax valuation, current availability and the credit constraint. Merely saying the student forgot discounting misses the economically important extension.

8.3.2 Negative saving and final consumption

With incomes 100 and 100, current tax 20, interest 10% and initial consumption 95, calculate saving and final consumption. Saving is -15 and final consumption is $100 - 16.5 = 83.5$. With future tax 22 and the same initial consumption, saving is 5 and final consumption is $78 + 5.5 = 83.5$. Equal consumption pairs use different private contracts. Without borrowing, the first choice is inadmissible and the second remains feasible. Check limits before comparing scores.

8.3.3 Changing the temporal weight

Use the $\beta = 0.5$ extension. Explain why higher benchmark initial consumption does not contradict financing neutrality. The announced preference differs from the baseline, but within this extension preferences match across calendars and feasible sets remain identical. Approximately 113.94 applies to both. Comparing only with baseline 89.52 mixes the preference dimension with the tax-calendar effect. The expected saving difference remains 20 even though both levels of saving can be negative.

8.3.4 A fund in two states

Calculate the external contribution with future-tax revenue 11 and 33. Repayment is always 22: the fund contributes 11 in the first state and receives 11 in the second. Its signed expected contribution is zero under equal probabilities. Expected household utility need not equal utility of expected consumption because the logarithm is concave. A correct answer distinguishes a public financing identity from private risk valuation and recognises that a zero expected fiscal transfer does not eliminate risk.

8.3.5 Partial compensation through transfers

Under intermediate weight, current-tax transfer is zero and future-tax transfer is 13.75. Recipient consumption is 50 versus $28 + 1.10(13.75) = 43.125$. The positive transfer difference does not fully offset the final tax of 22; its initial compensation ratio is 0.6875. Explain how partial compensation can be optimal with a transfer constraint. A departure from one is not necessarily a calculation mistake or neglect of the future obligation.

8.3.6 Two means with an incomplete sample

Suppose three people finish current tax but only one finishes future tax. The report can count complete choices for progress, but should not treat the difference in those means as four complete-block changes. Use only the block with both choices and declare the denominator. Explain why global revelation waits for full completion. Matched exclusions avoid comparing different sets of people, while gating prevents the results from informing choices that are still pending.

8.3.7 Interest and the income profile

Compare present-value resources for $(60, 140)$ and $(140, 60)$ under current tax at 10%. They are $40 + 140/1.10 \simeq 167.27$ and $120 + 60/1.10 \simeq 174.55$. Total income is 200 in both, but earlier income has greater present value. Then replace current tax with future tax within each profile and recover the same respective wealth. This exercise shows precisely which comparison neutrality applies to and prevents confusing total undiscounted income with intertemporal resources.

8.4 Data, privacy and inference limits

Exports preserve episode, calendar, incomes, role, consumption, balance, taxes, rates, utility, score and benchmark. Forecasts appear where relevant, and completed records are marked. Technical session and participant identifiers join repeated records but should never appear in projected personal tables. The aggregate report contains no names, student numbers or these identifiers. Guide captures use test sessions and impersonal labels; they are not evidence of real student responses.

Several records from one person share experience and characteristics. Fixed pairs also share a relationship, while rematched participants carry history into new relationships. Do not treat every exported row as independent without a justified analysis structure. Teaching use supports descriptive means and comparisons. Research requires a protocol, hypotheses, assignment and a plan addressing dependence. Passing bots and reconciled accounts verifies software behaviour, rather than replacing an experimental research design.

Do not collect extra demographic information by default. A written reflection can explain a budget without linking private entry codes to a projected table. Preserve session databases rather than resetting them to update screens. When sharing outcomes, use an aggregate summary stating condition, participants, decisions and complete matched blocks. These counts let another instructor understand what was observed and the denominator supporting an interpretation.

8.5 Common problems and an operating checklist

If a field rejects consumption, inspect the interval and decimal entry before assuming a payment failure. Choices must preserve positive final consumption in every announced state. If a pair waits, check that the first person submitted consumption and the second submitted a forecast. Closing a browser tab does not automatically supply a response. Individual baseline choices do not wait for another person; unexpected waiting suggests a generation institution or the wrong configuration.

Pending records after first-episode outcomes mean later active episodes remain. A dash for a mean indicates a zero denominator, not a mean of zero. If consumption differs from a benchmark, check comprehension, parameters and constraints before advancing a behavioural interpretation. Under risk, distinguish state-dependent realised points from the expected-utility benchmark. If Demo or Rooms shows another family, inspect the selected access before running the class.

Bibliography

- Barro, R. J. (1974). Are Government Bonds Net Wealth? *Journal of Political Economy*, 82(6), 1095–1117. <https://barro.scholars.harvard.edu/publications/are-government-bonds-net-wealth>.
- Cadsby, C. B., and Frank, M. (1991). Experimental Tests of Ricardian Equivalence. *Economic Inquiry*, 29(4), 645–664. <https://doi.org/10.1111/j.1465-7295.1991.tb00852.x>.
- Meissner, T., and Rostam-Afschar, D. (2017). Learning Ricardian Equivalence. *Journal of Economic Dynamics and Control*, 82, 273–288. <https://doi.org/10.1016/j.jedc.2017.07.004>. Author copy: <https://rostam-afschar.de/papers/JEDC2017.pdf>.
- Ricciuti, R., and Di Laurea, D. (2003). An Experimental Analysis of Two Departures from Ricardian Equivalence. *Economics Bulletin*, 8(11), 1–11. <https://flore.unifi.it/handle/2158/221023>.
- American Economic Association. JEL Classification System. <https://www.aeaweb.org/econlit/jelCodes.php?view=jel>.
- Behavioral and Experimental Economics Research Lab. Teaching game catalogue. <https://teaching.beer-lab.org/>.