



Instructor guide for the inflation expectations lab

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 inflation expectations lab 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 Design, objectives, and preparation	1
1.1 Access through the updated laboratory website	1
1.2 What the class will learn	5
1.3 Family summary	5
1.4 Institutional access and four rooms	6
1.5 Preparation before students arrive	7
2 Participant and instructor experience	8
2.1 Welcome, instructions, and comprehension	8
2.2 Simultaneous forecasts and available history	8
2.3 Outcomes, rewards, and closing	13
2.4 Create, monitor, and reveal	13
3 Economic foundations and worked examples	17
3.1 Feedback and the no-shock reference	17
3.2 A fully reconstructed economy	18
3.3 Learning, stability, and reference rules	18
3.4 What the incentive measures and does not measure	19
4 Reading the report and experimental evidence	20
4.1 Errors, rewards, and denominators	20
4.2 Histogram, trajectory, and persistence	20
4.3 What the literature contributes	21
5 Extensions and their predictions	24
5.1 No feedback: forecasting an exogenous process	24
5.2 Negative feedback: reversing collective influence	24
5.3 Strong feedback: fixed points and instability	25
5.4 Higher uncertainty: noise and disagreement	25
5.5 Public projection: a verifiable common reference	25
5.6 Private signals: information about the disturbance	26
5.7 Announced target: anticipating a regime change	27

5.8	Two horizons: forecasting and settling the correct dates	28
5.9	Inflation and output: a triangular system	28
5.10	Limited history: what information remains available	30
6	Class plan, assessment, and data	32
6.1	A 75-minute class	32
6.2	Assessable exercises and answer guidance	32
6.3	Export, privacy, and reproduction	33
6.4	Troubleshooting and operational closing	33

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
2.1	Welcome: read and confirm the rules before starting.	9
2.2	Current instructions for the basic edition: initial resources, decision rules, and experimental units.	10
2.3	Current comprehension check with an incorrect answer, its specific explanation, and retained responses.	11
2.4	Decision: enter the current forecast in the stated unit.	12
2.5	Waiting: keep the page open while the economy completes its decisions.	13
2.6	Individual outcome: distinguish the error, reward, and economy outcome.	14
2.7	Closing: keep the page open for discussion.	14
2.8	Creation: choose a configuration and a compatible participant count.	15
2.9	Initial report: progress is visible and outcomes remain hidden.	16
3.1	Instruction example: check the rule before calculating a reward.	18
4.1	Report: read each horizon using its own denominator.	21
4.2	Histogram: interpret the tails without removing large errors.	22
4.3	Trajectory: compare inflation, mean forecasts, and target.	22
4.4	Benchmarks: compare accuracy against the same realizations.	23
5.1	Public projection: observe a common number without an accuracy guarantee.	26
5.2	Private signal: distinguish an imperfect clue from the exact shock.	27
5.3	Two horizons: verify the target dates of both fields.	29
5.4	Two variables: enter forecasts for the same date.	30

Design, objectives, and preparation

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.

ión	Laboratorio de
ión	Laboratorio de expectativas de inflación / Inflation Expectations Lab — BEER Lab Seleccione una variante de este juego. Select a version of this game.
os: grupo fijo	
os: reasignación	
referencia	
in realimentación	
realimentación negativa	
realimentación intensa	
mayor incertidumbre	
proyección pública	
señal privada	
cambio anunciado de objetivo	
Los horizontes	
inflación y producto	
historia limitada	
tion	
fixed group	
reassignment	
reference	
o feedback	
egative feedback	
trong feedback	
igher uncertainty	
ublic projection	
ivate signal	
nnounced target change	
wo horizons	
flation and output	
imited history	

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

Rooms

Current rooms:

[ES] Expectativas: básica
[EN] Expectations: basic
[ES] Expectativas: extendida
[EN] Expectations: extended

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

Room: [EN] Expectations: basic

Create a new session

Session Config

[EN] One interaction

Number of participants

Complete groups of 6 participants. Use a multiple of 6; minimum 6.

Create

[Configure session](#) 

App sequence

expectations_lab_game (20 rounds)

Six simultaneous inflation forecasts feed back into a transparent macro outcome.

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 What the class will learn

The expectations lab examines a relationship that is often hidden in macroeconomic data: people try to anticipate an outcome, while their expectations also help produce it. Each economy contains six analysts. Everyone submits a point forecast of inflation; the program aggregates those responses, adds a shock, and calculates realized inflation. Each person then observes forecast accuracy and earns a reward. This sequence turns the abstract idea of endogenous expectations into a chain that the class can reconstruct numerically.

By the end, students should distinguish a forecast from a policy decision, explain the difference between individual error and collective disagreement, and construct a counterfactual while holding the shock constant. They should also recognize whether a history belongs to the same economy and whether the people composing that economy have changed. An intermediate course can add derivation of the no-shock reference, calculations for an informative signal, and comparisons between current and ahead horizons. These objectives can be assessed through short explanations and worked problems; success does not require a particular classroom outcome.

The outcome is automatic, with no student choosing an interest rate. It therefore provides a comparison that complements the monetary policy game already available. There are no asset purchases, deposit creation, or credit negotiations. The observable object is a forecast and its relationship with inflation rather than a financial transaction. One extension introduces output to teach two related variables while preserving that distinction.

Theoretical foundation

An endogenous outcome depends on the choices of the people trying to anticipate it. A shock is a perturbation assigned by the program and unknown when a decision is made. Separating these components is essential: high realized inflation alone does not show that the group forecast high inflation.

1.3 Family summary

Purpose

Understand feedback, learning, and inflation forecast accuracy.

Experimental unit	Economy of six analysts; economy-period and individual forecast observations.
Basic edition	One interaction; ten periods with a fixed group; ten periods with reassignment across economies.
Extended edition	Twenty-period reference and ten useful extensions, with fixed groups.
Decision	Current point forecast; an additional horizon or output when requested by the configuration.
Payoff	Up to 100 ECU per settled forecast: $100 \exp[-(f - \pi)^2/4]$. Components are added within each period.
Suggested duration	30–40 minutes for one interaction; 45–60 for ten periods; 60–75 for twenty.
Participants	Multiples of six; at least twelve for reassignment. Everyone forecasts.
Topics	inflation expectations; forecasting; learning; feedback; uncertainty; macroeconomic dynamics.
JEL	A22; C92; D84; E31; E37; E71.

The suggested duration includes instructions and discussion. Students familiar with the platform may decide quickly, but compressing discussion removes the main learning opportunity. A six-person class can use a fixed economy; twelve people create two economies and allow meaningful reassignment. With larger classes, settle participation at a multiple of six before creating the session. Do not add a person without a function or replace missing forecasts with instructor responses.

1.4 Institutional access and four rooms

Begin at <https://beer-lab.org/> and choose the teaching games link. At <https://teaching.beer-lab.org/>, find the inflation expectations lab card and open the edition to be used. The catalog supports discovery through topics and JEL codes. Check the title, language, and edition before distributing a room, because a participant link cannot change treatment once a session has started.

Family domain:

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

Instructor access:

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

Room list:

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

The four stable student routes are:

- Spanish basic:
https://expectationslab.beer-lab.org/room/expectativas_es.
- English basic:
https://expectationslab.beer-lab.org/room/expectations_en.
- Spanish extended:
https://expectationslab.beer-lab.org/room/expectativas_extendido.

- English extended:
https://expectationslab.beer-lab.org/room/expectations_extended.

Demo must show 28 configurations from this family: three basic and eleven extended per language. **Rooms** must show exactly four rooms. A basic room offers three configurations and initially selects one interaction in its own language. An extended room offers eleven, beginning with the reference. If another family appears, suspend session creation and contact laboratory staff. The filtered list supports selection, while the independent database separates sessions and exports.

A room assigns individual links to arriving participants; the instructor uses private access for creation and monitoring. An individual link retains a person's technical identity within the session and must not be shared between students. The instructor report has a session-specific administrative route, accessible from session start links. That route does not replace a room and is not distributed to students.

1.5 Preparation before students arrive

First decide which comparison you want to teach. One interaction is useful for reconstructing the equation and distinguishing a forecast, a mean, and an outcome. Ten periods support discussion of adaptation. Twenty periods are preferable for an extension that changes information or adds a second horizon. Avoid chaining many conditions into one class: prior experience changes behavior and consumes time intended for interpretation. For a clear comparison, select the reference and one extension, keep the same horizon, and prepare a question that can be answered with a report indicator.

Check connectivity, devices, and attendance. Create a synthetic practice session and verify entry, closing, and the report before inviting students into the room. This establishes functionality but does not replace a pedagogical pilot with students. Prepare a slide defining the symbols and keep results hidden. For the local instance, instructor access is <http://localhost:8087/demo>; it works only on this computer. Screenshots in this guide come from that implementation, and the public family routes require deployment.

For the instructor

Announce only the necessary rules: six forecasts, an automatic outcome, and a reward for accuracy. Reserve terms such as “coordination”, “anchoring”, and “instability” for the later discussion. Do not suggest a number students should forecast to generate the pattern you hope to teach.

Participant and instructor experience

2.1 Welcome, instructions, and comprehension

The first screen shows institutional logos, an introduction to the laboratory, and general rules. Students confirm that they have read them before starting. They remain silent, make private decisions, and keep the page open. When they need help, they contact staff without communicating forecasts to other people. This acceptance does not collect personal characteristics or serve as a demographic questionnaire; it organizes participation in a controlled activity.

Instructions define the units. Inflation and the output gap are expressed in percentage points; rewards are expressed in ECU. For example, an increase in inflation from 2 to 3 is a one-percentage-point increase in the rate. It does not mean that the price level rose by one ECU or that inflation was automatically multiplied by a payoff factor. This distinction prevents confusion between a macroeconomic magnitude and the activity's incentive.

The rule displays the coefficient and target for the selected configuration. The program provides the shock's equally likely support without revealing the current draw before a decision. It also explains whether the same people remain together or are reassigned. Extensions add only information needed to act: a private signal, public projection, target schedule, second horizon, output, or a history limit. The general example uses reference parameters and states that the session-specific rule appears above; instructors should emphasize that rule when selecting another condition.

Three questions check outcome timing, the aggregate entering the equation, and the reward criterion. For the first, the correct answer is the end of the same period for the current forecast. The second identifies the mean of all six responses, including when that mean has a zero coefficient. The third identifies the distance between forecast and outcome. Someone who makes a mistake rereads and corrects it; comprehension mistakes do not cost ECU. Private assistance can clarify the procedure without recommending a forecast.

2.2 Simultaneous forecasts and available history

The decision screen shows the period, current target, coefficient, and previous inflation. Students enter a number between -10 and 20 , with at most one decimal place. This interval constrains the input; the automatic outcome can lie outside it under strong feedback. The system rejects out-of-range responses, nonfinite values, and excessive decimal precision. Instructors can demonstrate an input error during preparation, but class time should not be spent discovering bounds through trial and error when the rules already provide them.

From period two onward, students see the history of their assigned economy and their own forecasting history. With fixed groups, both describe experience with the same people. Under



Period 1 / 1

Forecasting activity

The Behavioral and Experimental Economics Research Lab in the Department of Economics at Tecnológico de Monterrey is a space for learning about economic decisions through controlled activities.

Laboratory rules

- Remain silent. Make individual, private decisions; do not talk or communicate with others.
- Follow the instructions. Do not refresh, close, or leave the page during the activity.
- If you need help, raise your hand and ask staff privately.

☐ I have read and accept these rules.

Start

Figure 2.1: Welcome: read and confirm the rules before starting.



Period 1 / 1

Instructions

Resources and unit of account

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

This session does not convert ECU into real money.

Inflation and the output gap are expressed in percentage points. Rewards are expressed in ECU.

At the start of each period, six people simultaneously forecast the inflation realized at the end of that same period. The current shock is not observed before deciding.

The same six people remain in the economy throughout the activity.

Inflation is calculated using this rule:

$$\pi = 2.00 + 0.60 \times (\text{Mean forecast} - 2.00) + \varepsilon$$

Each of these shock values has equal probability: -0.50 / 0.00 / 0.50

Each settled forecast earns $100 \times \exp(-\text{error}^2/4)$ ECU. Error = forecast – observed value. An exact forecast earns 100 ECU; a 2-point error earns 36.79 ECU. Rewards are added.

Rule example

If the target is 2, the mean forecast is 4, the coefficient is 0.6, and the shock is 0, inflation is 3.2. A forecast of 4 has an error of 0.8 and earns 85.21 ECU. This example uses those parameters; your session rule appears above.

Continue

Powered by [oTree](#)

Figure 2.2: Current instructions for the basic edition: initial resources, decision rules, and experimental units.



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



Period 1 / 1

Check the rules

When is inflation for your current forecast observed?

Only in the following period ▾

Your answer: Only in the following period. The correct answer is At the end of the same period. Review the timing, the economy rule, and the reward for accuracy.

Which value enters the inflation rule?

The mean of the six current forecasts ▾

What determines a forecast reward?

The distance between the forecast and realized value ▾

Continue

Powered by QTree

Figure 2.3: Current comprehension check with an incorrect answer, its specific explanation, and retained responses.

reassignment, macroeconomic history belongs to a continuing economy identifier, whereas personal history follows the student. People change without restarting the series. This explicit teaching choice allows the class to ask whether accuracy changes when the people generating the next outcome differ. The comparison should not be described as “the same partner” versus “another partner”: each unit contains six analysts.

When the form contains two horizons, the first field concerns inflation realized in the current period. The second concerns the next period. Both are submitted together, and only the first enters the current equation. When the form contains inflation and output, both fields concern the current period and earn separate rewards. The output gap accepts values from -10 to 10 . Read the full labels; two fields that look similar need not have the same target date.

Confirmation leads to a neutral waiting page. Six responses are needed to calculate the outcome, and the economy is released automatically when they are available. The program does not reveal partial responses or replace a slower person. Under reassignment, the entire session is synchronized before the next period so that every economy’s history is complete. This additional wait preserves chronology: no one uses an outcome as information before it exists.

The screenshot shows a web interface for the BEER Lab. At the top left are the logos for 'Tecnológico de Monterrey' and 'BEER'. At the top right, it says 'Period 1 / 1'. The main heading is 'Enter your forecasts'. Below this, there are three input fields with their current values: 'Current target' with '2.00', 'Coefficient' with '0.60', and 'Previous inflation' with '2.00'. Below these fields, a instruction reads: 'Use values between -10 and 20 , with at most one decimal place.' Then, there is a label 'Current-period inflation forecast' followed by an empty text input box. At the bottom left is a blue 'Continue' button. The entire form is enclosed in a light blue border.

Tecnológico de Monterrey | BEER

Period 1 / 1

Enter your forecasts

Current target	Coefficient
2.00	0.60

Previous inflation

2.00

Use values between -10 and 20 , with at most one decimal place.

Current-period inflation forecast

Continue

Figure 2.4: Decision: enter the current forecast in the stated unit.



2.3 Outcomes, rewards, and closing

The outcome screen separates inflation, the shock, and, when permitted by the condition, mean forecasts and dispersion. Each person sees their response, signed error, current reward, and cumulative reward. A positive error means the forecast exceeded inflation; a negative error means it was below inflation. Since the reward depends on the squared difference, equally large errors with opposite signs earn the same reward. Preserving signs in the report nevertheless identifies directional biases hidden by average rewards.

With two horizons, from period two onward the screen displays the earlier response targeting the current date and its lagged reward. This does not evaluate the current forecast twice. With inflation and output, the output response and error appear alongside their payoff component. Components are summed and rounded to ECU cents once per period; precise scores remain available for analysis. A manual sum of already rounded displayed numbers can therefore differ by one cent without indicating a mechanism error.

The final screen explicitly states that the activity has ended, summarizes the total, and asks students to wait for instructions. It provides no button leading to an empty page. Students can inspect their payoff history while the instructor prepares the debriefing. Under two horizons, no prediction is requested for a date after closing: the last ahead response was submitted in the penultimate period and is settled in the last. No reward remains pending for an outcome in a nonexistent session.

What participants will see

Each page has one main title and buttons in the session language. To check the experience, open a complete synthetic session in the language you will use; inspecting only a welcome screenshot cannot reveal a mistranslated outcome label or second-horizon field.

2.4 Create, monitor, and reveal

The instructor signs into the panel and selects a room, configuration, and participant count. Keep the monitor and report in private tabs. The monitor helps locate an operational blockage; the report provides aggregates for projection. They serve different purposes. Do not project a monitor containing individual links or codes, or announce which person remains on a comprehension page. A persistent wait is resolved by locating the pending screen and helping privately.

Opening the report initially shows progress and expected denominators only. Outcomes remain

Realized inflation	2.56
Inflation shock	0.50
Mean forecast	2.10
Forecast dispersion	0.34

Forecast 1.60	Signed error -0.96
Current forecast reward 79.42 ECU	Total period reward 79.42 ECU
Cumulative reward 79.42 ECU	

Figure 2.6: Individual outcome: distinguish the error, reward, and economy outcome.



Period 1 / 1

The activity has ended

Your decisions have been recorded. Wait for the instructor's directions and keep this page open.

Cumulative reward: **79.42 ECU**

Period	Total period reward
1	79.42 ECU

Figure 2.7: Closing: keep the page open for discussion.

hidden until an explicit action, and reveal is enabled only when all economies have completed their active periods. Completing the first period of a twenty-period session is not the same as finishing the activity. A hide button restores a display without outcomes; reloading also restores the hidden state. Before projection, confirm the treatment, modality, and number of economies, because two administrative tabs may belong to different sessions.

Checkpoint

Before discussion, check completed active periods, the correct language, no personal identifiers, and denominators matching the comparison. In a one-interaction session, the report must count that interaction even if the administrator opens it from the application's final technical period.

Create session

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

Session Config

[ES] Una interacción



Number of participants

6

Must be a multiple of 6

Create

[Configure session](#)

App sequence

expectations_lab_game (20 rounds)

Six simultaneous inflation forecasts feed back into a transparent macro outcome.

Figure 2.8: Creation: choose a configuration and a compatible participant count.

Instructor report: inflation expectations

Reference · Fixed group

Participants

6

Economies

1

Completed economy-periods

0 / 1

Pending economy-periods

1

Recorded current forecasts

0

Aggregates use completed economy-periods only. Each economy-period contains six current forecasts. Ahead forecasts count only when their target period is settled.

Reveal becomes available when all economies have completed the active periods.

Reveal aggregate results

Figure 2.9: Initial report: progress is visible and outcomes remain hidden.

Economic foundations and worked examples

3.1 Feedback and the no-shock reference

Let f_{it} be person i 's current forecast for period t , and let $\bar{f}_t$ denote the mean of the six forecasts. The announced target is π_t^* , the shock is ε_t , and the feedback coefficient is γ . The basic equation is

$$\pi_t = \pi_t^* + \gamma(\bar{f}_t - \pi_t^*) + \varepsilon_t.$$

A one-point increase in the mean changes inflation by γ points, holding everything else constant. By contrast, increasing just one forecast by one point changes the mean by $1/6$ and inflation by $\gamma/6$. Confusing these experiments exaggerates individual influence sixfold. With $\gamma = 0.6$, changing one person's response from 2 to 3 raises inflation by 0.1 if the other five responses and shock remain unchanged.

Without a shock, when everyone forecasts the same value f , a fulfilled forecast satisfies $f = \pi^* + \gamma(f - \pi^*)$. If $\gamma \neq 1$, the solution is $f = \pi^*$. The reference value is 2. This calculation identifies a fixed point; it does not show that a class will reach it or that everyone uses the same learning rule. Strong feedback preserves this fixed point while potentially amplifying deviations. The existence of a solution and stability of a trajectory are separate questions.

To examine the individual incentive, hold the sum S of the other five forecasts fixed and remove the shock. The person's error is

$$e_i = (1 - \gamma/6)f_i - (1 - \gamma)\pi^* - (\gamma/6)S.$$

Because the reward decreases with e_i^2 , an interior best response makes the error zero. Hence,

$$f_i = \frac{(1 - \gamma)\pi^* + (\gamma/6)S}{1 - \gamma/6}.$$

With $\gamma = 0.6$, a target of 2, and five others forecasting 4, $S = 20$, yielding $2.8/0.9 \simeq 3.111$. Since the form accepts one decimal place, 3.1 approximates this response. Exactly copying the others' mean is not always the best response because one's own response also enters the outcome.

Attention

This solution assumes no shocks and known forecasts from others. During the activity, responses are simultaneous and current responses from others are unobserved. Do not use the formula as though students know S , or describe every deviation as irrational without considering their information.

3.2 A fully reconstructed economy

Consider six responses: 1, 2, 2, 3, 4, and 6. Their sum is 18 and their mean is 3. With target 2, coefficient 0.6, and shock 0.5, inflation is $2 + 0.6(3 - 2) + 0.5 = 3.1$. The person responding 1 has error -2.1 and earns $100 \exp(-4.41/4) \simeq 33.20$ ECU. The person responding 3 has error -0.1 and earns approximately 99.75 ECU. The person responding 6 has error 2.9 and earns approximately 12.21 ECU. No one is rewarded for ranking first: all scores depend continuously on distance.

Forecast dispersion is calculated relative to the mean of 3. The six deviations are $-2, -1, -1, 0, 1$, and 3 . Their squared values sum to 16. Variance is therefore $16/6 \simeq 2.667$, and the standard deviation is approximately 1.633 points. The denominator is six because this indicator describes the six responses composing that economy, rather than estimating a hypothetical population with an unbiased sample statistic. Group dispersion is neither the error of the person responding 3 nor the shock of 0.5.

Now hold the shock at 0.5 and replace all six responses with 2. Inflation becomes 2.5; every error is -0.5 , and the reward is approximately 93.94 ECU. Dispersion falls to zero, but error does not because the shock remains. This comparison explains why perfect agreement need not produce perfect accuracy. It also shows that the original responses changed the outcome by 0.6 points independently of the perturbation.

For a second counterfactual, restore the original responses and set $\gamma = 0$. Inflation is 2.5 even though the mean is 3. The outcome matches the case where everyone forecast 2 under positive feedback, but for a different reason. In one case responses changed; in the other their influence was removed. Identical outcomes do not identify mechanisms without observing choices and parameters.

Rule example

If the target is 2, the mean forecast is 4, the coefficient is 0.6, and the shock is 0, inflation is 3.2. A forecast of 4 has an error of 0.8 and earns 85.21 ECU. This example uses those parameters; your session rule appears above.

Figure 3.1: Instruction example: check the rule before calculating a reward.

3.3 Learning, stability, and reference rules

A simple adaptive rule copies the last observed inflation. If everyone uses it and shocks are removed, the deviation $d_t = \pi_t - \pi^*$ satisfies $d_t = \gamma d_{t-1}$. With $\gamma = 0.6$ and initial inflation of 4, subsequent inflation values are 3.2, 2.72, and 2.432. Distance shrinks at every step. With $\gamma = -0.6$, the values are 0.8, 2.72, and 1.568: they alternate around the target with decreasing amplitude. With $\gamma = 1.2$, they become 4.4, 4.88, and 5.456: deviations grow. These trajectories follow a particular updating rule, not an unconditional prediction for every class.

The extrapolative benchmark adds the last inflation change to the last inflation value. If history ends with 3 followed by 4, it forecasts 5. This rule can summarize a trend, but that trend may reflect temporary shocks. When many people extrapolate under positive feedback, their reaction can reinforce the movement they are trying to anticipate. The adaptive benchmark would forecast 4 in this example. The report evaluates both rules against the same outcomes and restricts them to the permitted interval, comparing alternatives a participant could actually enter.

The first period uses two initial inflation values of 2. These are not observations generated by previous student choices. They are announced initial conditions that make both benchmarks equal

to 2. When the target changes, mechanically copying history can lag even though the schedule is known. A temporary difference between inflation and the new target may reflect learning, historical dependence, or uncertainty about others' responses. The report describes that difference but does not directly observe a mental rule.

Disagreement and deviation of the mean have different roles. If half the group forecasts 0 and half forecasts 4, the mean is 2 and the no-shock outcome is 2 despite high dispersion. If everyone forecasts 4, dispersion is zero and the no-shock outcome is 3.2. The first case combines individual disagreement with no aggregate departure from the reference; the second combines collective agreement with a departure. Drawing both cases is useful before equating stability with uniformity.

3.4 What the incentive measures and does not measure

The reward function has a maximum of 100 when error is zero. A one-point error earns approximately 77.88 ECU, and a two-point error approximately 36.79. The decline is nonlinear: reducing error from two points to one raises the reward by about 41.09 ECU; reducing it from one to zero adds about 22.12. A low reward is not a loss from an initial endowment. No endowment is deducted and no additional fine is imposed for a distant forecast; rewards are always nonnegative.

The incentive collects a point prediction rather than a complete distribution. Under uncertainty, maximizing $E[\exp(-(f - \pi)^2/4)]$ is not generally equivalent to minimizing $E[(f - \pi)^2]$. Minimizing squared error yields the conditional mean when the variable is exogenous. The exponential reward weights distances differently, and in this game the realization also depends on one's own forecast. A response must therefore not automatically be interpreted as the mean of subjective beliefs.

This limitation does not remove the teaching value: everyone has an explicit accuracy incentive, and the report preserves untransformed errors. MSE, bias, and mean reward can be compared without treating them as identical. If conditions generate different numbers of settled forecasts, compare rewards per forecast and errors by horizon. Across twenty periods, the reference settles twenty forecasts per person; two horizons settles thirty-nine; inflation and output settles forty. Corresponding maximum totals are 2000, 3900, and 4000 ECU before any real-money conversion, which this activity does not specify.

Feedback script

"First identify the date we forecast and the outcome realized. Then calculate the untransformed error. Only afterward should we examine the reward. Higher ECU may result from greater accuracy or from more settled forecasts; we need the denominator to distinguish them."

Reading the report and experimental evidence

4.1 Errors, rewards, and denominators

The report separates progress from aggregate outcomes. An economy-period is complete when all six responses are available and its outcome has been calculated. Recorded decisions track progress but cannot enter a mean independently when their inflation outcome does not yet exist. With two economies and ten periods, the maximum is twenty economy-periods and 120 current forecasts. If one economy is missing in the last period, nineteen units are complete and 114 forecasts are usable in aggregates; reveal remains unavailable.

For errors e_j in one variable and horizon, mean squared error is $\text{MSE} = \sum_j e_j^2 / N$, and bias is $\bar{e} = \sum_j e_j / N$. The first preserves magnitude; the second preserves direction and can cancel opposite errors. Errors of -1 and 1 produce zero bias and MSE of 1. Errors of 0 and 0 produce zero for both. The first pair is less accurate despite identical bias. Mean reward transforms each error before averaging; it is not the transformation of mean error.

The decomposition $\text{MSE} = \text{Var}(e) + \bar{e}^2$ uses population variance over these observations. Changes in MSE may reflect lower error dispersion, lower bias, or both. Forecast dispersion within an economy is a different indicator, calculated before subtracting the common outcome. Do not mix the standard deviation of six forecasts with variance in errors accumulated across dates having different inflation outcomes.

The two-horizon report separates current and ahead inflation rows. A complete session contains twenty responses per person in the first row and nineteen in the second. An output row exists only when that variable is elicited. A dash means absent observations or insufficient variation, never an estimated zero. In a classroom exercise, ask students to write the numerator, denominator, and unit of each figure before making comparisons. This often exposes horizon confusion faster than a discussion of total rewards.

4.2 Histogram, trajectory, and persistence

The histogram distributes signed current inflation errors. Its intervals include both tails: below -2 and at least 2 . Large errors are not removed because realizations lie outside the response interval. Bars sum to the number of settled current forecasts. Concentration around zero indicates descriptive accuracy; a positive tail indicates frequent overprediction. Sample size and shock support must accompany this description. Six observations from one economy in one period do not estimate a stable population distribution.

Settled forecasts

Variable and horizon	N	Mean squared error	Mean signed error	Mean reward per forecast
Current inflation	120	0.29	-0.14	93.26 ECU
One-period-ahead inflation	114	0.21	-0.22	95.01 ECU

Mean distance from target

0.41

Inflation lag correlation

-0.18

N = 19

Correlation uses consecutive pairs within the same economy; it requires at least three and variation in both series. It does not compare people or establish a mental rule. Dispersion is the population standard deviation of the six forecasts.

Figure 4.1: Report: read each horizon using its own denominator.

The trajectory overlays realized inflation, mean forecasts, and the target on one scale. Each point averages complete economies at that date, and the table retains their count. A stable mean between two economies can conceal one rising while the other falls. If comparison requires heterogeneous trajectories, use the private economy-level export and draw separate graphs. Do not attribute a path to an average student when it was first averaged across groups.

Report persistence is a descriptive correlation between inflation and its lag, based on consecutive pairs within each economy. It never links one economy's last period to another's first. At least three pairs and variation in both series are required. Positive correlation may reflect learning, shocks, or collective responses; it does not independently identify an adaptive rule. Under reassignment it still describes the economy process even though human composition changed.

Adaptive and extrapolative benchmarks compare accuracy against the same outcomes. Their MSE uses one reference prediction per economy-period, whereas participant MSE uses six. Both are comparable averages of errors, but their observation counts differ. A benchmark with lower MSE does not establish that students should have known it or that anyone used it. Ask which information the rule used and which counterfactual comparison it represents.

4.3 What the literature contributes

Pfajfar and Žakelj (2014) study inflation expectation formation in a New Keynesian experimental environment. Their work motivates attention to heterogeneous forecasting rules and learning rather than assuming one universal forecasting approach. The BEER family retains incentivized forecasts interacting with an economy but uses a transparent linear equation. It does not reproduce the study's parameters, duration, or estimated percentages of participant types.

Current inflation error distribution

N = 120. Signed-error intervals, in percentage points. All finite errors are included, including those outside the permitted forecast input interval.





Figure 4.2: Histogram: interpret the tails without removing large errors.

Economy evolution

Realized inflation Mean forecast Current target



Each point is the mean across completed economies in that period; the table preserves denominators. Target, mean forecast, and inflation share one scale.

Figure 4.3: Trajectory: compare inflation, mean forecasts, and target.

Forecast benchmarks

Forecast	Mean squared error
Last observed inflation	0.38
Last inflation plus last change	1.12

Benchmarks are calculated before observing the current shock. They are restricted to the response interval -10 to 20 . Both initial values are 2 in the first period. Errors use the same set of economy-periods; they do not attribute strategies to participants.

Figure 4.4: Benchmarks: compare accuracy against the same realizations.

The review by [Hommes \(2021\)](#) explains how positive feedback and coordination of forecasting rules can generate persistent fluctuations. This synthesis motivates comparisons of coefficient signs and strength; it does not guarantee that twenty classroom periods will generate the published paths. [Rholes and Petersen \(2021\)](#) compare communication of point projections and uncertainty. Their evidence shows that adding information need not improve every indicator. This activity implements an automatic point projection and a simplified private signal; neither is claimed to replicate their density treatments.

[Assenza et al. \(2021\)](#) connect self-organization of expectations with macroeconomic outcomes and policy. The BEER inflation-output extension takes the motivation of observing two variables but uses its own triangular system without a monetary instrument. This adaptation teaches relationships without leaving implicit simultaneity unresolved or repeating the central bank game. The references serve different purposes: paradigm, synthesis, communication, and relationships between variables.

Feedback script

“The literature provides mechanisms and comparisons; our session provides observations from this class. First describe what happened. Then identify a compatible mechanism and the additional information needed to distinguish it from another explanation.”

Extensions and their predictions

The extended edition contains a twenty-period reference and ten extensions. Each changes a mechanism with an identifiable teaching question. Keep groups fixed to interpret history and horizons; compare an extension with a reference of the same length, not with one interaction. The number reflects usefulness and adequate specification rather than a ceiling imposed by the thematic plan. Simplified magnitudes and rules are BEER adaptations grounded in the cited literature's mechanisms and are identified as such.

5.1 No feedback: forecasting an exogenous process

The coefficient changes from 0.6 to zero, preserving the target and shock distribution. Inflation is $2 + \varepsilon_t$ regardless of the six forecasts. The teaching question is what changes when outcomes stop responding to expectations. Uncertainty and the need to decide remain; strategic aggregate influence disappears. The mean remains a defined variable in the equation even though its term has a zero coefficient.

Without a shock, 2 is exactly accurate. With shocks, a short history can fluctuate even if everyone responds 2. Compare MSE, dispersion, and the visual relationship between mean forecasts and inflation. A mean far from 2 while inflation stays near the target is consistent with an exogenous environment; it does not show different student ability. The review by [Hommes \(2021\)](#) motivates separating learning in exogenous and endogenous processes, but this exact coefficient is a teaching choice.

Exercise: six people forecast 8 and the shock is -0.5 . Inflation is 1.5. If everyone changes to 2 with the same shock, inflation remains 1.5. Ask which counterfactual was possible in the reference but is absent here. The expected answer identifies the mean-forecast channel while holding the perturbation fixed. Do not describe this condition as an absence of expectations: forecasts are observed but do not influence the outcome.

5.2 Negative feedback: reversing collective influence

With $\gamma = -0.6$, increasing the mean lowers inflation when the shock is held fixed. The question is whether learning in an environment responding in the opposite direction produces different behavior. The no-shock reference remains 2; reversing the sign does not itself change the fixed point. If everyone copies previous inflation, deviations alternate in sign and contract. That prediction depends on the specified rule and should be assessed descriptively rather than treated as inevitable.

Compare successive target deviations, dispersion, and accuracy. A zigzag around 2 may arise from learning or alternating shocks. To separate them, reconstruct at least one date using the observed

mean and shock. A mean of 4 with no shock yields inflation of 0.8. If the next mean were 0.8, subsequent no-shock inflation would be 2.72. This sequence explains how copying the last outcome can induce oscillation.

The feedback synthesis in [Hommes \(2021\)](#) supports comparing signs. The BEER equation, shock support, and horizon are specific to this activity. Ask whether a negative relationship between mean forecasts and inflation indicates that people tried to be wrong. The correct answer is no: the economy rule determines the sign, while everyone remains rewarded for individual accuracy.

5.3 Strong feedback: fixed points and instability

This extension uses $\gamma = 1.2$. The target remains 2, but departures of the mean can produce larger departures in the outcome. The teaching question distinguishes equilibrium existence from stability under an updating rule. If everyone forecasts 2 and there is no shock, inflation is 2. If everyone copies previous inflation above the target, deviations grow under the adaptive rule. Neither statement requires an explosive classroom path.

Inflation is never clipped. A mean of 20 and shock of 0.5 yields $2 + 1.2(20 - 2) + 0.5 = 24.1$, outside the forecast interval. The form limits actions rather than the generated variable. In this state an exactly accurate response may be unavailable. Explain this property before discussing precision because high MSE can reflect the input constraint rather than only slow learning.

Inspect graph scales, tail errors, and target distance. The review by [Hommes \(2021\)](#) motivates amplification; 1.2 and the response bounds are explicit adaptations. Ask how the economy would change if the program clipped 24.1 to 20. Answers should recognize a new rule, a nonlinearity, and different rewards. Do not treat clipping as a merely visual correction.

5.4 Higher uncertainty: noise and disagreement

The coefficient stays at $\gamma = 0.6$, while shock support expands to $\{-1.5, 0, 1.5\}$. The shock mean remains zero. Its variance rises from $(0.25 + 0 + 0.25)/3 = 1/6$ to $(2.25 + 0 + 2.25)/3 = 1.5$, a ninefold increase. The question is how accuracy changes with a larger perturbation while the expectation influence mechanism remains unchanged.

If everyone forecasts 2, inflation is the target plus shock. With a shock of 1.5, every error is -1.5 and the reward is approximately 56.98 ECU. Under the reference shock of 0.5, it would be 93.94. Forecast dispersion is zero in both cases. This separates lower accuracy from greater disagreement. Do not present all inflation variation as evidence of unanchored expectations.

Compare MSE and dispersion, reconstructing shock and mean contributions. The expectations literature reviewed by [Hommes \(2021\)](#) and communication work by [Rholes and Petersen \(2021\)](#) motivate distinguishing uncertainty and information; the larger discrete support is a teaching adaptation. Ask whether ninefold shock variance requires observed MSE to be exactly nine times larger. It does not: responses and their endogenous influence may also change, and the sample is short. The ratio describes the assigned shock rather than a guaranteed empirical effect.

5.5 Public projection: a verifiable common reference

Before deciding, everyone observes $q_t = \pi_t^* + 0.6(\pi_{t-1} - \pi_t^*)$. This automatic projection uses previous inflation only. It does not observe the current shock, change the equation, or constitute a promise. The question is whether a common number changes disagreement and forecast

coordination. Matching the projection is observable; believing the institution is credible is an additional interpretation not directly elicited here.

If previous inflation is 4 and the target is 2, the projection is 3.2. If everyone copies it and there is no shock, new inflation is $2 + 0.6(3.2 - 2) = 2.72$. The projection does not perfectly predict the outcome generated by its own adoption. This example teaches a subtle point: a past-based number can guide actions whose aggregation changes a future outcome. Do not call the projection an infallible rational prediction.

Rholes and Petersen (2021) motivate studying communication content and format; this activity uses a simpler rule without a human authority. Compare response distance from q_t , dispersion, and MSE using the export and report. Ask what information the projection used and what it could not know. Students should identify the lag and absence of the current shock. Reduced disagreement alongside higher error would show why coordination and accuracy require separate evaluation.

The screenshot shows a web interface for a forecasting exercise. At the top, there are logos for 'Tecnológico de Monterrey' and 'BEER', and a 'Period 1 / 20' indicator. The main heading is 'Enter your forecasts'. Below this, there are four light blue boxes arranged in a 2x2 grid, each containing a label and a value: 'Current target' with '2.00', 'Coefficient' with '0.60', 'Previous inflation' with '2.00', and 'Public projection' with '2.00'. Below these boxes is a text instruction: 'Use values between -10 and 20, with at most one decimal place.' followed by a label 'Current-period inflation forecast' and an empty text input field. At the bottom left is a blue 'Continue' button.

Figure 5.1: Public projection: observe a common number without an accuracy guarantee.

5.6 Private signals: information about the disturbance

Each person observes $s_{it} = \varepsilon_t + \nu_{it}$, where private noise is $-0.5, 0, 0.5$ with equal probability and conditional independence across people. The common shock retains those same three values. The signal informs students about the current disturbance without revealing others' responses. The question is how a group turns heterogeneous information into forecasts that also affect outcomes.

This is a transparent adaptation of information in the expectations paradigm rather than a replication of a published information structure.

There are nine equally likely shock-noise pairs. A signal of 0.5 is compatible with $(\varepsilon, \nu) = (0, 0.5)$ and $(0.5, 0)$. Both have equal conditional probability, so $E[\varepsilon \mid s = 0.5] = 0.25$. A signal of 1 is compatible only with $(0.5, 0.5)$, identifying the shock. For signals $-1, -0.5, 0, 0.5, 1$, conditional means are $-0.5, -0.25, 0, 0.25, 0.5$. This table describes shock information; it does not automatically solve the optimal inflation forecast because others' responses must still be anticipated and the reward is exponential.

Compare dispersion, MSE, and response sensitivity to signals using private exports. Greater disagreement can reflect different information rather than poorer understanding. Ask whether six different signals imply six different economic shocks. They do not: everyone shares the same disturbance but observes private noise. The projected report does not display person-level signals; analyze that relationship only in private data or hypothetical examples.

Tecnológico de Monterrey | BEER | Period 1 / 20

Enter your forecasts

Current target 2.00	Coefficient 0.60
Previous inflation 2.00	Your private signal 0.50

Use values between -10 and 20 , with at most one decimal place.

Current-period inflation forecast

Continue

Figure 5.2: Private signal: distinguish an imperfect clue from the exact shock.

5.7 Announced target: anticipating a regime change

The target is 2 through period 5 and 4 from period 6 onward. Instructions announce the schedule from the start, and each decision displays the current target. There is no secret policy shock or student changing the target. The teaching question is whether earlier history remains influential

when a parameter change is known in advance. The no-shock reference is 2 before the change and 4 afterward.

If everyone still forecasts 2 in period 6 and there is no shock, inflation is $4 + 0.6(2 - 4) = 2.8$. If everyone forecasts 4, the outcome is 4. Forecasting 3 produces 3.4. These differences do not arise from uncertainty about the schedule: they may reflect expectations of others' responses, historical dependence, or misunderstanding. Check comprehension before interpreting delay as a psychological bias.

Inspect target distance, mean forecasts, and dispersion before and after the change. Communication work by [Rholes and Petersen \(2021\)](#) motivates attention to announcement timing; the deterministic BEER shift does not reproduce their policy. Ask why copying last inflation and using the new target yield different forecasts in period 6. Request a counterfactual holding the shock fixed. Do not compare raw inflation levels as though the target were unchanged: the parameter changed, and the relevant measure must acknowledge it.

5.8 Two horizons: forecasting and settling the correct dates

In every period except the last, students submit f_{it} for current inflation and h_{it} for next inflation. Only f_{it} enters the outcome at t . The reward for h_{it} is calculated against π_{t+1} and paid at $t + 1$. The question is how accuracy changes when information is further from the target date. Forecasting and communication research, including [Rholes and Petersen \(2021\)](#), motivates observing horizons; this adaptation keeps ahead forecasts outside the equation to make chronology unambiguous.

Example: in period 1 a person submits $f_{i1} = 2$ and $h_{i1} = 3$. If $\pi_1 = 2.5$, the current reward is 93.94 ECU. The second forecast is not evaluated against 2.5. If $\pi_2 = 3.5$, it is subsequently settled with error -0.5 and reward 93.94. Period 2 also contains a new current forecast earning its own reward against 3.5. The two reward components in that period come from decisions submitted at different dates.

The final decision does not request a forecast for period 21. Across twenty periods, each person submits twenty current and nineteen ahead forecasts. Inspect errors and rewards by horizon with their denominators, and distinguish submission dates from evaluation dates. Ask what would change if h_{i1} were evaluated against π_1 . The answer identifies a different task rather than an innocuous label change. That error destroys intertemporal interpretation even when payoff totals seem plausible.

5.9 Inflation and output: a triangular system

Everyone submits an inflation forecast f_{it} and an output-gap forecast g_{it} for the same date. First calculate $x_t = 0.5\bar{g}_t + d_t$; then $\pi_t = 2 + 0.6(\bar{f}_t - 2) + 0.4x_t + \varepsilon_t$. Demand and inflation shocks are independent, each taking $-0.5, 0, 0.5$ with equal probability. The question is how one endogenous variable enters another without leaving an unresolved simultaneous system. The gap is a reduced measure relative to reference output, not observed employment or GDP recorded by real firms.

If $\bar{g}_t = 2$, $d_t = 0.5$, $\bar{f}_t = 4$, and $\varepsilon_t = 0.5$, first obtain $x_t = 1.5$ and then $\pi_t = 4.3$. A person forecasting output of 2 has error 0.5; someone forecasting inflation of 4 has error -0.3 . Their rewards are approximately 93.94 and 97.78 ECU, respectively, added as two components. Without shocks and with responses $g = 0$, $f = 2$, both forecasts are fulfilled.

[Assenza et al. \(2021\)](#) motivate studying expectations and joint macroeconomic variables; the BEER triangular system, coefficients, and rewards are original adaptations. Compare MSE by variable and explain calculation order. Ask how much inflation changes if mean output forecasts rise by



Period 1 / 20

Enter your forecasts

Current target 2.00	Coefficient 0.60
Previous inflation 2.00	


Use values between -10 and 20, with at most one decimal place.

Current-period inflation forecast

Next-period inflation forecast

Continue

one point, holding everything else fixed: output rises by 0.5 and inflation by 0.2. The answer is not 0.4, which links inflation to realized output rather than directly to the output forecast.



Tecnológico de Monterrey | **BEER** Behavioral Economics Experiment Research Laboratory Period 1 / 20

Enter your forecasts

Current target

2.00

Coefficient

0.60

Previous inflation

2.00

Use values between -10 and 20, with at most one decimal place.

Current-period inflation forecast

For the output gap, use values between -10 and 10, with at most one decimal place.

Current-period output-gap forecast

Continue

Figure 5.4: Two variables: enter forecasts for the same date.

5.10 Limited history: what information remains available

The screen retains only the previous three periods, and individual outcomes omit mean forecasts and dispersion. Rewards and the equation remain unchanged. The question is how available information affects forecast revision. Instructors retain full histories in the report and export. The three-period window is a specific teaching choice, motivated by the literature's interest in learning and information rather than attributed as an exact published treatment.

In period 7, students can consult dates 4, 5, and 6. Personal history shows those three responses and experienced outcomes; economy history shows its three realizations. An event at date 2 is no

longer displayed, although a student may remember it. The restriction affects the screen without removing human memory or preventing all mental reconstruction. Do not attribute an outcome to pure forgetting without measuring memory.

Inspect dispersion, MSE, and benchmark performance. Keep the horizon constant because more periods provide different learning opportunities. Ask whether removing the others' mean changes the value entering the equation. It does not: the change concerns information revealed afterward, not aggregation or economic influence. If accuracy is similar, describe that result and consider that three recent observations may have been sufficient. A small difference does not show historical information lacks value in every environment.

Class plan, assessment, and data

6.1 A 75-minute class

Allocate ten minutes to access, units, and sequence; five to comprehension; twenty to periods; ten to privately reviewing the report; twenty to discussion; and ten to an individual closing exercise. The total is 75 minutes. Adjust session length before creation to actual time available, rather than cutting an active two-horizon session with unsettled rewards. For a short class, one interaction allows thirteen minutes for instructions and comprehension, three for decisions and waiting, eight for the report, and six for assessment.

During decisions, restrict verbal intervention to procedure. After closing, ask what determined observed inflation, which part depended on mean forecasts, and which part reflected the shock. Reveal metrics before the trajectory. Request description without causal explanation: level, variation, disagreement, and accuracy. Then reconstruct one date and compare a counterfactual. This sequence prevents a striking graph from displacing the equation's basic accounting.

Compare reference and extension using sessions of the same length. If the same students complete both, disclose order and accumulated experience. Alternating order across groups improves teaching comparisons without eliminating all contamination. Statistical significance is not required for a successful class. An outcome near the reference supports discussion of convergence; a distant outcome supports reconstruction of mechanisms and limits. Assess learning through explanation and calculation rather than production of an expected curve.

For the instructor

Choose one main question: expectation influence, feedback sign, uncertainty, information, or forecast timing. One carefully reconstructed comparison teaches more than running ten extensions without interpretation time. The complete guide supports selection; it does not prescribe running everything in one class.

6.2 Assessable exercises and answer guidance

Exercise 1: individual influence. Five forecasts are 2 and the sixth person changes from 2 to 8. There is no shock and $\gamma = 0.6$. The mean rises from 2 to 3, so inflation changes from 2 to 2.6. A six-point individual change produces a 0.6-point economic change. Multiplying six directly by 0.6 confuses a personal response with the mean. Assess whether students keep the other five decisions fixed and define the change's unit.

Exercise 2: accuracy and agreement. Compare six forecasts of 4 with three forecasts of 0 and three of 4, without a shock. In the first case inflation is 3.2, everyone has error 0.8, and dispersion

is zero. In the second inflation is 2, errors are -2 and 2 , MSE is 4, and dispersion is 2. Complete agreement can coexist with bias, while a mean at the target can coexist with large errors. Answers must separate three objects: outcomes, errors, and dispersion.

Exercise 3: dates. An ahead response of 4 was submitted in period 3. Inflation was 3 then and 4.5 in the next period. The correct reward is evaluated against 4.5, yielding error -0.5 and approximately 93.94 ECU, settled in period 4. Evaluation against 3 would yield 77.88 ECU and measure another task. Request a timeline showing submission, realization, and payment; award chronology credit even when there is a small arithmetic error.

Exercise 4: inference. Two classes have different MSE, but one played after hearing the other's discussion. The difference is descriptive and does not identify a clean treatment effect. Causal attribution requires comparable assignment, information, and experience. A complete answer proposes a better comparison and recognizes dependence among people in one economy rather than treating each forecast as an independent economy.

6.3 Export, privacy, and reproduction

The private panel exports session, economy, period, language, edition, treatment, matching, and seed. Individual records include position, forecasts, target date for the second horizon, submission time, shown information, precise score components, rounded payoff, completion status, and economic outcome. Participant identifiers link personal trajectories within the private file; they never appear in the projected report. Keep any association with names outside public discussion.

The reconstruction unit is session–economy–period. Grouping by economy alone mixes sessions reusing group numbers. Grouping only by person confuses personal and macroeconomic histories under reassignment. To verify a row, reconstruct the six-response mean, inspect target and shocks, and apply the equation. Evaluate each horizon's error against the correct date. Preserve input data before creating derived variables and document any exclusion of incomplete periods.

Technical tests and screenshots in this delivery use synthetic participants. They demonstrate functionality rather than behavior in a student sample. Research use requires an appropriate protocol, hypotheses, and applicable institutional procedures. This guide is a teaching tool and does not replace that specification. Ordinary classes should protect links, exports, and privacy under laboratory practices without collecting additional demographics.

6.4 Troubleshooting and operational closing

If a page waits indefinitely, check the monitor and confirm six connected people in every economy. Someone still on comprehension or facing a validation error has not submitted a forecast. Help privately. Under reassignment, a finished economy may wait for another before the next period; this synchronization is part of the design and prevents incomplete histories. Do not force an outcome to release the session.

If a link reports that the user no longer exists, it probably belongs to an earlier database or session. Create a new session and distribute current links; inventing a user cannot recover a trajectory. Restarting the local process preserves its database and does not require deleting sessions. If the application needs updating, coordinate restart before class and recheck Demo, Rooms, and languages. Changing configurations does not transform already collected data into another treatment.

For report dashes, check denominators and variation. One interaction does not provide enough lag observations for persistence. If reveal is disabled, inspect all active periods rather than the latest

outcome seen by one person. For text or screenshots in another language, verify that the correct room was opened. Local practice credentials are admin/admin; published instances use assigned institutional credentials.

Checkpoint

At closing, confirm all decisions, reveal aggregates only, export privately, preserve the configuration and treatment order, and record which explanation students could reconstruct. Keep technical incidents with the session so they are not confused with economic outcomes.

A useful final assessment asks for a reconstructed equation, an interpreted graph, and a counterfactual. Students should explain what changed, what remained constant, and why rewards alone do not summarize the mechanism. This combination links operation, theory, and evidence without attributing conclusions to a class that its data cannot support.

Bibliography

- Pfajfar, D., and Žakelj, B. (2014). Experimental evidence on inflation expectation formation. *Journal of Economic Dynamics and Control*, 44, 147–168. <https://doi.org/10.1016/j.jedc.2014.04.012>.
- Hommes, C. (2021). Behavioral and Experimental Macroeconomics and Policy Analysis: A Complex Systems Approach. *Journal of Economic Literature*, 59(1), 149–219. <https://doi.org/10.1257/jel.20191434>.
- Rholes, R., and Petersen, L. (2021). Should central banks communicate uncertainty in their projections? *Journal of Economic Behavior & Organization*, 183, 320–341. <https://doi.org/10.1016/j.jebo.2020.11.013>.
- Assenza, T., Heemeijer, P., Hommes, C., and Massaro, D. (2021). Managing self-organization of expectations through monetary policy: A macro experiment. *Journal of Monetary Economics*, 117, 170–186. <https://doi.org/10.1016/j.jmoneco.2019.12.005>.
- American Economic Association. JEL classification, accessed October 4, 2026. <https://www.aeaweb.org/econlit/jelCodes.php?view=jel>.
- oTree. Documentación oficial de formularios, rondas and bots. <https://otree.readthedocs.io/>. Local implementation verified with oTree 6.0.15.