



Instructor Guide: Deterrence, Reputation, and Credibility

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 helps instructors prepare, implement, and discuss the BEER Lab game of deterrence, reputation, and credibility. It is written for instructors with different levels of familiarity with experimental economics and international relations.

The document offers three reading paths. The quick guide supports session launch. Operational chapters explain participant screens and the instructor report. Conceptual chapters develop payoffs, predictions, variants, and evidence for the classroom discussion.

Common rules are explained once. Each variant then states only what changes from the basic game: information, incentives, capacity, or history. Worked examples and exercises help interpret outcomes without attributing a single observed action to an unmeasured mechanism.

Author

This guide was prepared by Stanislao 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 interface. Spanish configurations implement the same logic. Never project individual links or connect choices to names or seats.

Contents

How to use this guide	i
1 Quick implementation guide	1
1.1 Access through the updated laboratory website	1
1.2 Activity summary	4
1.3 Preparing in ten minutes	5
1.4 A feasible class schedule	6
2 What the game represents	7
2.1 A crisis with unequal information	7
2.2 Order, information, and unit of analysis	7
2.3 One encounter worked from beginning to end	8
3 Reference predictions and theoretical limits	9
3.1 Backward induction without additional measures	9
3.2 Messages, signals, and selection	9
3.3 Other preferences, uncertainty, and learning	10
4 Screens and classroom facilitation	11
4.1 From the portal to individual links	11
4.2 Turns, waiting, and completion	12
5 Reading the instructor report	18
5.1 Reveal and progress	18
5.2 Three levels of interpretation	19
6 Basic modes and extended treatments	20
6.1 One encounter, fixed partner, and reassignment	20
6.2 No prior message	20
6.3 Known resistance cost	21
6.4 Backing-down cost	21
6.5 Contingent deposit	21
6.6 Verified history	22
6.7 History with imperfect observation	22
6.8 Limited defensive capacity	23
6.9 Material preparation	23
7 Class plan, exercises, and solutions	24
7.1 Before, during, and after the activity	24
7.2 Exercise 1: the resistance threshold	24
7.3 Exercise 2: challenging under a subjective belief	25
7.4 Exercise 3: threats, deposits, and selection	25

7.5	Exercise 4: Bayesian updating with noisy history	25
7.6	Exercise 5: ex ante versus ex post	25
8	Data, privacy, and common problems	26
8.1	Export and statistical unit	26
8.2	Troubleshooting	26
8.3	Final operational checklist	27
9	Comparison and inference in class	28
9.1	Counterfactuals and assignment	28
9.2	What is and is not measured	28
9.3	Welfare, conflict, and unmeasured outcomes	29
10	Evidence and references	30
10.1	How to use the literature	30

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.	3
1.4	Rooms contains the four rooms for this family: basic and extended editions in Spanish and English.	3
1.5	Session-creation form for the basic English room. Check the selected configuration and participant requirement before creating the session.	4
4.1	The opening page requires the general laboratory rules to be acknowledged.	11
4.2	Current instructions for the basic edition: initial resources, decision rules, and experimental units.	13
4.3	Current comprehension check with an incorrect answer, its specific explanation, and retained responses.	14
4.4	Each person sees the assigned role before the strategic sequence begins.	14
4.5	A sends a nonbinding message and remains free to choose differently later.	15
4.6	Waiting synchronizes choices without revealing the partner's response prematurely.	15
4.7	B observes the permitted information and chooses whether to challenge.	16
4.8	After a challenge, A compares the private cost with the visible winning probability.	16
4.9	The individual result shows the observed path and own payoff without personal data.	17
4.10	The final screen confirms the total and ends the activity explicitly.	17
5.1	The report shows progress while keeping outcome distributions hidden.	18
5.2	After reveal, rates, denominators, outcome paths, and encounter trends become visible.	19

Quick implementation guide

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.

Reputación y credibilidad — Un encuentro	Disuasión, reputación y credibilidad / Deterrence, Reputation, and Credibility — BEER Lab Seleccione una variante de este juego. Select a version of this game.
Reputación y credibilidad — Diez encuentros con pareja fija	
Reputación y credibilidad — Diez encuentros con nueva pareja	
Reputation, and Credibility — One encounter	
Reputation, and Credibility — Ten encounters with a fixed partner	
Reputation, and Credibility — Ten encounters with partner reassignment	
Reputación y credibilidad — Sin mensaje previo	
Reputación y credibilidad — Costo de resistencia conocido	
Reputación y credibilidad — Costo de retractación	
Reputación y credibilidad — Depósito contingente	
Reputación y credibilidad — Historial verificable	
Reputación y credibilidad — Historial con observación imperfecta	
Reputación y credibilidad — Capacidad defensiva limitada	
Reputación y credibilidad — Preparación material	
Reputation, and Credibility — No prior message	
Reputation, and Credibility — Known resistance cost	
Reputation, and Credibility — Backing-down cost	
Reputation, and Credibility — Contingent deposit	
Reputation, and Credibility — Verified history	
Reputation, and Credibility — History with imperfect observation	
Reputation, and Credibility — Limited defensive capacity	
Reputation, and Credibility — Material preparation	

Figure 1.3: Demo contains only the configurations of this game family, in both languages and editions.

Rooms

Current rooms:

[ES] Disuasión, reputación y credibilidad
[EN] Deterrence, Reputation, and Credibility
[ES] Disuasión, reputación y credibilidad — extendido
[EN] Deterrence, Reputation, and Credibility — extended

Figure 1.4: Rooms contains the four rooms for this family: basic and extended editions in Spanish and English.

Room: [EN] Deterrence, Reputation, and Credibility

Create a new session

Session Config

[EN] Deterrence, Reputation, and Credibility — One encounter

Number of participants

Complete groups of 2 participants. Use a multiple of 2; minimum 2.

Create

[Configure session](#) 

App sequence

deterrence_game (10 rounds)

Deterrence, Reputation, and Credibility / Disuasión, reputación y credibilidad.

Figure 1.5: 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 Activity summary

Purpose	Examine when a threat of resistance deters a challenge and how information, capacity, commitment, and reputation change its credibility.
Experimental unit	A pair consisting of A, who holds the initial position, and B, who may challenge it. The number of participants must be even; reassignment needs at least four people.
Basic edition	One encounter; ten encounters with a fixed partner; ten encounters with partner reassignment. Roles remain constant.

Extended edition	No message, known cost, backing-down cost, contingent deposit, verified history, history observed with error, limited capacity, and material preparation.
Decision	A announces an intention and sometimes takes a prior measure. B chooses whether to challenge. After a challenge, A yields or resists.
Payoff	Without a challenge, A receives 100 and B 60 ECU before any preparation cost. Yielding gives A 60 and B 100 before additional costs. Resistance leads to one winner draw and the specified conflict costs.
Suggested duration	35–45 minutes for one encounter; 60–80 minutes for ten encounters and discussion.
Participants	An even number; at least four for reassignment.
Topics	deterrence; threat credibility; reputation; private information; commitment; international conflict.
JEL	A22; C72; C92; D74; F51; F52.

1.3 Preparing in ten minutes

Confirm the number of people first. Two participants suffice for a one-encounter session, whereas reassignment needs at least four and is more useful with several pairs. Every participant receives a distinct link. Do not give the same link to several students. Project a waiting screen or the instructor report, never a private choice page. Allow roughly five minutes for the comprehension check, ten to fifteen minutes for the first interaction, and fifteen to thirty minutes for comparison and discussion. If your aim is reputation, use ten encounters and allow time for students to learn from earlier actions.

Begin at <https://beer-lab.org/>, follow the teaching catalog at <https://teaching.beer-lab.org/>, and find the card “Deterrence, Reputation, and Credibility”. On the proposed family domain, <https://deterrence.beer-lab.org/>, the instructor demonstration is <https://deterrence.beer-lab.org/demo> and the room list is <https://deterrence.beer-lab.org/rooms>. The basic English room is https://deterrence.beer-lab.org/room/deterrence_en; the Spanish room is https://deterrence.beer-lab.org/room/disuasion_es. The extended rooms are https://deterrence.beer-lab.org/room/deterrence_extended and https://deterrence.beer-lab.org/room/disuasion_extendido. These are intended product routes. Check with the lab whether the domain has been deployed before using a public address in class.

The room lists configurations for its language and edition. An instructor creates a session, checks its participant count, keeps the private report link, and distributes one individual link to each person. Individual links depend on that session’s database. If the database is reset, old links stop working; create a new session and distribute its new links. The family’s **Demo** and **Rooms** views should contain only deterrence configurations and four rooms. If other games appear, do not launch the session and ask the lab coordinator to check the instance filter. A room for the basic edition should begin with a basic configuration in the same language.

For the instructor

Choose the mechanism you want students to compare before class. To separate words from incentives, pair the nonbinding message with a backing-down cost or deposit. To teach information, compare hidden and known resistance costs. To teach reputation, use repeated

reassignment and ask what a new B can infer from A's previous action.

1.4 A feasible class schedule

A one-encounter session can take 35–45 minutes: five minutes for access, eight for rules and comprehension, ten for decisions, ten for revealing the report, and ten for discussion. Ten encounters may require 60–80 minutes. Explain that points are measured in ECU and calculated by the platform. Do not imply that ECU are cash unless a conversion rule has been approved and communicated. Remind the class that the two roles see different information. Silence during the activity keeps the message inside the platform as the only permitted communication channel.

What the game represents

2.1 A crisis with unequal information

A holds a valuable initial position, while B receives less. B may challenge to improve their outcome. A challenge does not by itself determine the result: A sees it and then chooses to yield or resist. If A yields, B gets the preferred outcome without conflict. If A resists, both face a probabilistic contest and costs. A threat to resist can thus alter B's decision even when fighting never takes place. The primary observable sign of deterrence is no challenge, but that observation has meaning only in light of what B believed A would do if challenged. A threat that is never tested is not direct evidence of follow-through.

The design is intentionally abstract. A and B are strategic positions, not real countries and not moral labels. The game does not represent civilians, treaties, geography, domestic institutions, or every dimension of war. Students are not assumed to speak for real states. The model reduces a complicated crisis to a decision tree so that the instructor can isolate the initial status quo, private resistance cost, fighting capacity, and credible commitment. This simplification also makes the inferential limits visible: a choice reveals an action, whereas its underlying belief or motive remains unobserved unless separately elicited.

Each participant has a 30 ECU reserve already incorporated in every final payoff. With no challenge, A receives 100 and B 60. When B challenges and A yields, A receives 60 and B 100 before any special charge. If A resists, A pays a private cost of either 5 or 25 and B pays 30. One draw gives A a baseline one-half probability of winning. If A wins, A gets 130 minus their resistance cost and B gets zero. If B wins, A gets 30 minus their cost and B gets 100. These numbers already include the reserve. Do not add 30 again while explaining the results.

2.2 Order, information, and unit of analysis

The platform assigns a fixed role at the start. A receives a private cost, 5 or 25 with equal probability, and keeps that cost across repeated encounters. B knows the distribution and payoff rules. Except in the known-cost treatment, B does not see the current partner's cost before choosing. A may send "I will resist", "I will yield", or no message. The no-message treatment omits this stage. A visible prior measure, when applicable, precedes the message. B then chooses challenge or no challenge. Only a challenge activates A's response choice. A conflict uses exactly one saved winner draw for that pair and encounter.

The pair is the experimental unit of an encounter because actions are connected and payoffs are jointly determined. A person who participates in ten encounters contributes repeated, correlated observations; ten choices from one student are not ten independent students. Roles remain fixed so that a switch of strategic position does not become confused with the treatment. With a fixed

partner, the same two people interact again. With reassignment, each A receives a new B and each B a new A; the immediately previous pairing is avoided. In a small class a pair may reappear after more than one encounter. Mention that limitation when discussing how a reputation might form.

What participants will see

The opening page identifies BEER Lab and asks participants to accept the general rules. The instructions summarize common payoffs, reserve, uncertainty, private cost, and the session-specific rule. Three comprehension questions precede any strategic choice. Each person sees the role and encounter number; A sees the private cost, while B sees only information allowed by the treatment. Choice pages require a valid selection. Waiting pages do not reveal the other person's response prematurely.

2.3 One encounter worked from beginning to end

Suppose A has cost 5, announces resistance, and B challenges. A resists. Initially A has a 50% chance of winning. If the draw favors A, A receives the reserve of 30 plus a victory value of 100 minus the cost of 5, or 125 ECU. B receives a reserve of 30, no victory value, and pays a conflict cost of 30, or zero. If the draw favors B, A receives 30 minus 5, or 25, and B receives 30 plus 100 minus 30, or 100. Before seeing the draw, A's expected payoff from resistance is half of 125 plus half of 25, or 75. B's expected payoff in that branch is half of zero plus half of 100, or 50.

Repeat with A's cost 25. A would obtain 105 after a win or 5 after a loss, so the expected payoff from resistance is 55. Yielding gives 60. A high-cost person therefore prefers to yield under purely monetary preferences, while the low-cost person compares 75 with 60 and prefers to resist. B does not know which type is present in the basic game. B must consider the probability of actual resistance rather than taking the message literally. If both types announce resistance because they would like B to stay out, the message will not separate them. The observed response after a challenge is more informative than the unchecked claim, although it is visible only when B challenges.

Theoretical foundation

A threat is credible when carrying it out is attractive once the threatened contingency arrives. An intention can be sincere when stated and still cease to be optimal after a challenge. Analyze the tree backward: begin with A's response, then study B's challenge choice, and only then evaluate the value of an earlier statement or commitment device.

Reference predictions and theoretical limits

3.1 Backward induction without additional measures

Let c be A's resistance cost and p A's probability of winning. Resistance yields $30 + 100p - c$ in expectation, while yielding gives 60. Under expected-payoff maximization A resists when $30 + 100p - c \geq 60$, equivalently when $100p - c \geq 30$. At $p = 0.5$ the cost threshold is 20. Cost 5 lies below it and cost 25 above it. If another parameter combination made A indifferent, a sharp choice prediction would need a tie-breaking assumption. No such tie occurs for the two baseline types in this implementation.

Let r be B's subjective probability that A will resist. If B challenges, the expected payoff is $(1 - r)100 + r[100(1 - p)]$: yielding gives B 100, while resistance gives B 100 times A's loss probability. This expression reduces to $100(1 - rp)$. Not challenging gives 60. B challenges if $100(1 - rp) > 60$, or $rp < 0.4$. With $p = 0.5$, deterrence requires B to expect resistance with probability at least 0.8. This is a theoretical threshold, not a belief measured by the application. The instructor can ask which belief would be compatible with a decision under these monetary assumptions, while making clear that the actual belief is unknown.

Suppose each cost type is initially equally likely and only the low-cost type resists. Then $r = 0.5$. B assigns one quarter probability to resistance combined with A's victory. B's expected payoff from challenging is 75, above 60. An uninformative nonbinding message therefore does not deter a purely monetary B in this benchmark. This prediction does not mean every student must challenge. Students may distrust the announced probability, dislike risky outcomes, place intrinsic value on avoiding conflict, make arithmetic errors, or regard a promise as a social commitment. A classroom dataset alone does not identify which of these explanations applies to a particular person.

3.2 Messages, signals, and selection

The baseline message is cheap: it costs no ECU to send and changing one's response later does not change the payoff. Signaling theory asks whether different types would send different messages when the receiver reacts. A cheap message can focus attention or coordinate expectations, but separation is not guaranteed. A high-cost A benefits when B refrains from challenge and thus has a reason to imitate the low-cost A by announcing resistance. Have students distinguish three counts: announcements of resistance, challenges following an announcement, and actual resistance among challenged announcers. The third count excludes announcements never tested. One cannot classify such announcements as fulfilled merely because there was no challenge.

Few challenges may make a threat look effective, but they also leave little evidence of how A would respond. This is selection into observing the contingency: A's response exists only when B challenges. Imagine one challenge out of ten encounters and resistance in that sole case. A reported resistance rate of 100% would rest on one opportunity. It would not show that all ten A participants would have resisted. The dashboard includes numerators and denominators to keep a striking percentage from being confused with abundant evidence. A separate belief elicitation could help explain why B did not challenge; this game intentionally does not impose that additional measurement.

3.3 Other preferences, uncertainty, and learning

The monetary model is a reference, not a complete description of motivation. A might resist to avoid appearing weak, retaliate against a challenge, or remain consistent with a stated intention. Another A might yield despite a monetary advantage from resistance because the conflict lottery is unpleasant or because B might receive zero. B might avoid a profitable challenge from fairness concerns or risk aversion. Each account would need more evidence. The choices made here cannot by themselves identify a particular emotion or belief, so feedback should use phrases such as "consistent with" rather than declaring the motive established.

Repetition adds learning. With a fixed partner, B sees whether the same A follows through, and present action can affect future encounters. With reassignment, the immediate partner changes. In the verified-history treatment a new B sees summarized previous responses, letting a reputation travel between pairs. In the noisy-history treatment a false observation can unfairly strengthen or weaken that reputation. Ask whether participants treat one observed action as definitive proof of a stable type, even though the action also depends on incentives, chance, and possibly error. The last encounter is especially useful: there is no later encounter in which a reputational benefit can be enjoyed.


Feedback script

"Before seeing the class data, what would a low-cost A do after a challenge? What would a high-cost A do? Now stand in B's position. If you do not know the cost, how valuable is a challenge under the 50% victory rule? Why might a forfeitable deposit tell you more than an unsupported sentence?" Wait for the reasoning, then display the observed rates together with their denominators.

Screens and classroom facilitation

4.1 From the portal to individual links

The institutional site leads to the teaching catalog and then to this game's descriptive card. On the family domain, the instructor can inspect configurations and launch a practice session through **Demo**. The **Rooms** page separates language and edition. A room is an entry point for students and leads to individual participant links; the report is a private session administration route. Never circulate the report link among students. The screens in this guide come from an isolated local instance of the finished game. They illustrate actual interface states; they are not evidence that the proposed public domain is already live.



 **BEER** BASIC · NONBINDING MESSAGE
Behavioral and Experimental Economics Research Lab
Tecnológico de Monterrey

Welcome to BEER Lab

The Behavioral and Experimental Economics Research Lab in the Department of Economics at Tecnológico de Monterrey is a place to learn how people make decisions through controlled activities.

In this activity, you will make private decisions with another person. Payments are expressed in ECU.

Laboratory rules

- Remain silent during the activity.
- Do not talk or communicate with other participants.
- Make decisions individually and privately.
- Follow the instructions on the screen.
- Do not refresh, close, or leave this page.
- If you need help, raise your hand or contact the staff privately.

☐ I have read the rules and am ready to start.

Start

Figure 4.1: The opening page requires the general laboratory rules to be acknowledged.

Once the session is created, ask participants to read the instructions and complete the comprehension questions without saying the answers aloud. Two questions check the 100 ECU payment in two

distinct paths; a third checks that B normally cannot see A's private cost. The known-cost variant is a stated later exception. A wrong answer leaves the participant on the form with a brief explanation. The application records the number of attempts for quality review, not as a personal grade. Do not request participant screenshots that expose private cost during play. If students need help, they should ask the instructor privately.

4.2 Turns, waiting, and completion

At each encounter both people see their role. A knows the cost and, where applicable, chooses a deposit or material preparation first. A may send a message. B waits until this prior choice is recorded. B's screen then shows the message, any visible measure, and only the cost or history permitted by the treatment. If B does not challenge, the result is 100/60 except for A's preparation cost. If B challenges, A's page shows the private cost and actual victory probability and requires yielding or resisting. The platform synchronizes the pair before displaying a result. Each person sees their own payoff and the final screen confirms their total.

A waiting page does not imply that data have been lost. Sequential decisions mean one person may wait longer while the partner reads. Ask students to keep the tab open. If one leaves, the dashboard will count unfinished encounters and the pair can remain stopped. Do not invent a response just to complete the classroom display. Continue discussion with completed pairs and state the denominator. In repeated sessions, check that participants have reached the next encounter before interpreting the most recent trend point. The final page is an explicit confirmation, avoiding an ambiguous technical end screen.

BASIC · NONBINDING MESSAGE

Activity rules

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.

In each encounter, one person is A and the other is B. A knows their private cost of resisting: 5 or 25 ECU. B knows the two costs are equally likely but normally does not know the cost of their partner.

A can send a nonbinding intention. B decides whether to challenge. If B does not challenge, A receives 100 and B receives 60 ECU before any preparation cost. If B challenges, A chooses between yielding and resisting. If A yields, A receives 60 and B receives 100 ECU before any additional cost or deposit in the variant.

If A resists, one winner is drawn. A initially has a 50% chance of winning. If A wins, A receives 130 minus their resistance cost and B receives 0; if B wins, A receives 30 minus their cost and B receives 100. The variant can change the probability or charge A for preparation.

Each payment already includes each person's 30 ECU reserve. Only information permitted by this variant is shown. Roles and A's private cost persist in repeated sessions.

Additional rule in this session

A's message does not bind their later action. Words and actions are recorded separately.

[Continue](#)

Con tecnología de [oTree](#)

Figure 4.2: 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.

Check the rules

With no challenge and no preparation, how much does A receive?

Your answer: 101. The correct answer is 100. Review the rules: A receives 100 without a challenge, B receives 100 if A yields, and A's private cost remains hidden in this session.

With a challenge and yielding, how much does B receive before any additional rule?

In this session, will B know A's private cost before deciding?

Continue

Con tecnología de [oTree](#)

Figure 4.3: Current comprehension check with an incorrect answer, its specific explanation, and retained responses.

ENCOUNTER 1 OF 1

Beginning of the encounter

Person B

One encounter


You are B. You will decide whether to challenge after seeing the information allowed in this session.

A's cost can be 5 or 25 ECU.

A's message does not bind their later action. Words and actions are recorded separately.

Continue

Figure 4.4: Each person sees the assigned role before the strategic sequence begins.



ENCOUNTER 1 · PERSON A

A's message

Which message will you send to B?

☐ If challenged, I will resist.


☐ If challenged, I will yield.

☐ I am sending no message.

A's message does not bind their later action. Words and actions are recorded separately.

[Continue](#)

Figure 4.5: A sends a nonbinding message and remains free to choose differently later.



ENCOUNTER 1 OF 1

Wait for the other person's decision

The encounter will continue when the other person has completed their choice. Keep this page open.

[Continue](#)

Figure 4.6: Waiting synchronizes choices without revealing the partner's response prematurely.

ENCOUNTER 1 · PERSON B

B's decision

Message received from A
If challenged, I will resist.

Do you challenge A?

☐ Do not challenge

☐ Challenge

Continue

Figure 4.7: B observes the permitted information and chooses whether to challenge.

ENCOUNTER 1 · PERSON A

A's response

B challenged you. What do you choose?

5 ECU
Your private cost of resisting

50%

A's message does not bind their later action. Words and actions are recorded separately.

☐ Yield

☐ Resist

If you resist, there is one draw for the winner. Check this session's costs before choosing.

Continue

Figure 4.8: After a challenge, A compares the private cost with the visible winning probability.




Figure 4.9: The individual result shows the observed path and own payoff without personal data.




Figure 4.10: The final screen confirms the total and ends the activity explicitly.

Reading the instructor report

5.1 Reveal and progress

The instructor report always shows the number of participants and expected, completed, and pending encounters. All aggregate outcomes remain hidden until the instructor presses “Reveal aggregates”. This avoids disclosing the class pattern on a projected monitor while decisions are still in progress. A hide button removes the aggregates from projection. The visible state persists in browser session storage for that tab and session; it does not alter any choice or database record. If another browser opens the report, check its visible state again rather than assuming it inherited the first browser’s choice.




Figure 5.1: The report shows progress while keeping outcome distributions hidden.

The challenge rate counts completed encounters with a challenge over all completed encounters. The resistance rate uses completed challenged encounters as its denominator. The threat follow-through rate counts announcements of resistance followed by actual resistance, among announcements of resistance that were challenged. If there is no denominator, the dashboard shows a dash instead of zero. A zero would mean an event had opportunities to occur but did not; a dash means no valid opportunities exist for that rate. State the unit, numerator, and denominator out loud before offering a behavioral interpretation.

Mean payments to A and B use all completed encounters. The path table separates no challenge, concession, and conflict; the counts sum to completed encounters. The private-cost table shows how many low- and high-cost A players were exposed, how many were challenged, and how many resisted. A’s histogram uses 20 ECU bins. A bar represents observed payment frequency, not anyone’s subjective victory probability. The trend table reports counts and means within each




Figure 5.2: After reveal, rates, denominators, outcome paths, and encounter trends become visible.

encounter. In an active session the latest round may contain only a few finished pairs; do not compare an incomplete row with a fully observed earlier row without saying so.

5.2 Three levels of interpretation

Begin with an observable fact, such as “seven of twelve completed encounters included a challenge”. Offer a compatible mechanism next: “B may have doubted the threat would be carried out”. Then state the limit: B’s beliefs were not elicited, so the mechanism is not established. A high conflict rate is not automatically a measure of aggression. It could reflect many challenges, a high resistance rate conditional on challenge, a different distribution of costs, or lottery outcomes. Ask students to reconstruct the decision path before discussing causation.

An instructor may launch separate sessions for different treatments and record their aggregates side by side. The dashboard summarizes one session; it does not pool sessions or estimate confidence intervals. Class composition, time, and order of presentation can change responses. A classroom comparison is effective for teaching mechanisms and creating hypotheses but is not automatically confirmatory evidence. For a formal analysis, export the data and specify the assignment unit, exclusions, repeated observations, and appropriate uncertainty estimates in advance. Never treat every encounter by one participant as independent evidence.

Checkpoint

Before discussing a percentage, identify its numerator, denominator, and experimental unit. A dash means the rate could not be calculated; it is not evidence of zero behavior. If encounters are pending, report how many are missing and avoid calling the display the final result for the entire class.

Basic modes and extended treatments

6.1 One encounter, fixed partner, and reassignment

One encounter is well suited to introducing backward induction. Each pair sees one sequence and cannot build a shared history. Ten encounters with a fixed partner allow statements, challenges, and responses to affect future interaction with the same person. A useful teaching contrast is between the response that seems best in the final encounter and a response that could preserve a reputation earlier. Do not ascribe every change to reputation: experience may improve arithmetic, reduce mistakes, or help the pair coordinate expectations. Ten encounters with reassignment keep roles constant but change the immediate counterpart. A private bilateral reputation does not automatically travel to a new partner unless a treatment supplies a visible record. The contrast helps separate an ongoing relationship from publicly portable reputation.

In any repeated mode, a cost type stays with A. Students can therefore learn about a stable feature rather than a freshly drawn type every round. This choice makes the history meaningful, while also producing correlation within person. Immediate prior partners are excluded from the next pairing, but a later rematch remains possible in a small group. If two people meet again after several rounds, their choices may reflect direct experience even without the on-screen record. The instructor should be precise about what “new partner” guarantees and avoid claiming that all ten partners are distinct. Ask students whether their strategy would change if the type were redrawn independently at each encounter.

6.2 No prior message

This treatment removes A’s initial statement. B knows the general game rules but receives no partner-specific announcement. It provides a reference for asking whether communication changes challenges through information, persuasion, social norms, or simple attention. In a monetary benchmark where both cost types would send the same cheap threat, deleting that message does not change B’s payoff-relevant information. If fewer challenges occur with messages, the pattern is compatible with words shaping expectations or perceived obligations, but the challenge rate alone does not distinguish those channels. Compare completed challenges and ask which exact decision changed. Silence here is imposed by the program; it must not be interpreted as a voluntary choice by A to remain quiet.

The no-message condition also clarifies a measurement issue. In the baseline, a follow-through rate is defined only among announced threats that faced a challenge. Here there is no announcement opportunity, so the metric should not be treated as zero. There may still be resistance after a challenge. Present these as separate measures: the absence of prior communication and A’s later response. If students describe resistance as “breaking silence”, remind them that the response is a

new action at a different node, not a retroactive message. Literature on signaling motivates the distinction, while the exact message wording is part of this classroom implementation.

6.3 Known resistance cost

B sees the current A's private cost before choosing whether to challenge. If both maximize expected ECU and trust the stated probabilities, low-cost A resists and high-cost A yields after a challenge. Facing cost 5, B expects resistance and earns 50 from a challenge, less than 60 without one; the monetary prediction is no challenge. Facing cost 25, B expects yielding and earns 100 from a challenge; the prediction is a challenge. This benchmark is sharper than in the private-cost case because B need not infer the type. Compare challenge frequencies by disclosed cost and then inspect actual responses. If B challenges a low-cost A, compatible explanations include arithmetic error, risk preference, or doubt that A will resist; the game does not identify which without additional belief elicitation.

Make students calculate both roles rather than memorizing "low cost deters". The key is not that cost 5 makes conflict painless; resistance still produces a lottery. The key is that A's expected payoff from resistance, 75, exceeds 60 from yielding. B anticipates this and values the challenge at only 50. A cost of 25 reverses A's preference because resistance is worth 55. The comparison teaches how a private parameter can alter the credibility of an identical verbal threat. Crisis bargaining research motivates hidden versus common information, but the classroom values 5 and 25 are an explicit adaptation, not figures taken from a particular international incident.

6.4 Backing-down cost

In this treatment A loses an additional 20 ECU only if A announced resistance and then yields after a challenge. There is no charge for sending the message itself and none when B stays out. For a high-cost A, yielding after such an announcement gives 40, while resistance retains an expected value of 55. The earlier threat may become credible because its fulfillment is now relatively attractive at the later decision node. A low-cost A already preferred resistance. B may challenge less often after "I will resist" if the penalty is understood. But A may announce yielding or no message to avoid exposure. Analyze announcement frequency, challenges by announcement, and follow-through when the threat was tested together.

The 20 ECU cost is an exogenous institutional penalty. No actual domestic audience observes a leader or chooses a sanction in this game. Tomz's audience-cost experiments motivate the broader idea that backing down can be politically costly; this classroom rule represents the incentive change without reproducing those experiments. A useful question is whether every penalty that makes resistance more likely is beneficial. A policy that reduces challenges might also make a crisis harder to resolve peacefully after a mistaken challenge. Compare both no-challenge and conflict rates, not just the frequency of successful threats. Mention that the model includes no welfare effect on uninvolved people, an important limitation for interpreting real crises.

6.5 Contingent deposit

A may post a 25 ECU deposit before sending a message. The deposit is returned if B does not challenge or if A resists, and forfeited if A yields after a challenge. It costs no points in the other paths, but it visibly changes the price of retreat. A high-cost A who deposited receives 35 from yielding, below the 55 expected from resistance; low-cost A also prefers resistance. A deposit can

therefore support a threat in a way a costless sentence cannot. If B sees a deposit, a monetary B may be more inclined to avoid challenge. Compare the deposit choice with the message, challenge, and final response, and ask whether B noticed that it was contingent on the path.

Deposit status is an endogenous choice. A may decline to deposit in order to keep flexibility, and willingness to deposit could correlate with private cost or an unmeasured taste for consistency. Consequently, a simple comparison of challenges against depositors and nondepositors does not isolate a causal deposit effect. Have students propose an alternative experiment that randomly assigns the opportunity or obligation to post a deposit while holding the monetary rule fixed. This introduces the difference between a choice that signals type and a device that directly changes incentives. It also shows that commitment can have two effects at once: it can make a response more likely and communicate information about the person choosing it.

6.6 Verified history

This treatment uses ten encounters with reassignment. Before B decides, B sees a summary of the current A's prior responses when challenged. An entry exists only if A actually faced a challenge; an encounter without challenge does not reveal whether A would have resisted. The record follows A across different partners, creating a minimal form of portable reputation. A may value resisting today because a future B might see the action and avoid a challenge. That future benefit can be larger early in the sequence than in the last encounter, when no further encounter remains in which to exploit the record.

Look for differences in challenges after observed resistance versus observed yielding, changes in resistance over time, and variation across private costs. These are descriptive patterns, not a proof that reputation itself caused a choice. Low-cost A players are already more likely to resist and may subsequently face fewer challenges after a strong history. Tingley and Walter's experimental work on repeated play motivates the question about strategic reputation building; the ten-encounter classroom format and payoff values are our own adaptation. Ask students whether one previous response is enough to infer type, and what other information would be needed to distinguish a stable preference from a response to one unusually attractive or unattractive situation.

6.7 History with imperfect observation

The system leaves actual actions and payoffs unchanged, but each response shown to a new B may be flipped independently with probability 20%. An A who resisted can appear to have yielded, and one who yielded can appear to have resisted. The private export records both the real history and the version shown. The treatment shows why reputations based on imperfect observation can be fragile or unfair. A rational observer should not treat one entry as infallible. A prior belief about resistance has to be combined with the reliability of the record. With a symmetric prior of 50% and an 80%-accurate binary observation, a "resisted" entry implies an 80% posterior in that simplified example. A different prior changes the posterior.

The application does not elicit subjective probabilities. We can describe how a Bayesian observer would update, but cannot claim that any particular participant did so. Compare challenges after favorable and unfavorable displayed histories while remembering that the visible history can be wrong. An observed effect of a false record on B's action is a pedagogical demonstration of informational vulnerability, not evidence that the person whose reputation suffered actually changed their behavior. Discuss why verified history and noisy history are not merely two different labels for the same condition: one changes the quality of what the partner sees while preserving the true response sequence.

6.8 Limited defensive capacity

If A resists, the probability of winning falls from 50% to 25%; private resistance costs remain unchanged. For cost 5, expected resistance pays $30 + 25 - 5 = 50$, below the 60 available from yielding. For cost 25, it pays 30, also below 60. Under the monetary benchmark both types therefore yield after a challenge, and B expects a payoff of 100 from challenging. A verbal threat has little material support in this environment unless preferences or beliefs differ from the benchmark. Compare conditional resistance with baseline, but distinguish the probability change from any change in wording: the text of the threat is the same.

If someone resists even with limited capacity, the accounting is not invalid. The decision may reveal nonmonetary concerns, a misunderstanding, or a belief that the win probability differs from the one displayed. It need not reflect irrationality. Ask students to compute A's threshold win probability for each cost and see exactly why 25% fails both. Then ask whether increasing B's conflict cost, holding A's capacity fixed, could deter B even if A were unlikely to resist. This extends the reasoning beyond the specific classroom parameter. The treatment is a compact representation of material capability; real conflicts involve uncertain, multidimensional capacity and changing alliances that this simple draw cannot capture.

6.9 Material preparation

A may pay 5 ECU at the beginning of an encounter to raise the chance of winning from 50% to 80%. B sees the preparation decision. The cost is paid even without a challenge, so prepared A receives 95 rather than 100 in a peaceful encounter. If prepared A resists, expected payoff is $30 + 80 - c - 5$: 100 for cost 5 and 80 for cost 25. Both exceed the 55 from yielding after preparation. B, expecting a prepared A to resist, would value a challenge at 20, below the 60 from no challenge. Preparation can thus sustain a threat through capability rather than through a penalty for retreat.

Preparation has a real opportunity cost. If B would not challenge anyway, A pays 5 and receives no strategic benefit in that encounter. The choice could still influence future partners in a repeated game, but this one-encounter treatment does not add such a reputational channel. Ask when a five-point expenditure is worthwhile given a probability that it prevents a challenge. Compare the preparation rate, challenge rate, and A's mean payoff. A low conflict rate among prepared pairs is not proof that preparation was free or welfare improving. It may also reflect selection: people who expected B to challenge may be more likely to prepare. The contrast with the deposit highlights two routes to credibility: changing one's strength in a fight and changing the cost of yielding.

Attention

These eight variants are classroom instruments inspired by mechanisms studied in international relations and game theory. Their exact figures, interface, and sequence are original teaching adaptations. None is an exact replication of a cited research protocol. The backing-down cost is automatic; no third player acts as a domestic audience or decides a social sanction.

Class plan, exercises, and solutions

7.1 Before, during, and after the activity

Before students arrive, verify that the local or deployed instance shows only this family, that the basic room opens in the correct language, and that you have enough distinct participant links. Choose one comparison that fits the class period: baseline versus known cost, or cheap message versus contingent deposit. If reputation is the focus, choose the ten-encounter history treatment and resist adding several other rules simultaneously. Prepare a board with three columns: monetary prediction, observed fact, and alternative explanation. This simple layout keeps a challenge rate from becoming an immediate judgment about students' personalities.

During entry, explain that A and B are abstract positions and that no one should disclose private cost or decisions. Ask participants to read the rules and answer the comprehension questions. Once choices begin, follow the completed and pending counters without opening private data. When the pairs have reached outcomes, reveal aggregates. Start with the three outcome paths, then break them into challenge frequency, conditional response, and lottery victory. If conflict occurs, remind students that winning does not prove resistance was optimal; the decision is evaluated before the draw. If there is no conflict, ask which untested threats could have deterred B.

At the end, return to the calculations for cost 5 and cost 25 and to B's threshold. The same statement can be credible for one A and not another because the incentive at the response node matters. Compare mechanisms that alter that incentive: a retreat penalty, a forfeitable deposit, or a higher chance of winning. In history treatments, discuss how a past action informs a new partner and why an error-prone observation warrants caution. Conclude by separating empirical association from causal effect and ask students to formulate a testable hypothesis for a later session.

7.2 Exercise 1: the resistance threshold

Ask: if A wins with probability 40% and has resistance cost 5, does expected ECU maximization favor resistance? The calculation is $30 + 40 - 5 = 65$, above the 60 from yielding, so yes. With cost 25, it is 45, so yielding is better. The general threshold is $p \geq (30 + c)/100$: 35% for cost 5 and 55% for cost 25. Ask where the 30 comes from. It is the difference between the 60 gained by yielding and the 30 reserve, not an additional war cost. A common error is to subtract the resistance cost twice because it appears in both the story and the formula.

7.3 Exercise 2: challenging under a subjective belief

Suppose B assigns probability $r = 0.6$ to resistance and thinks A wins conditional on resistance with probability $p = 0.5$. Challenging gives $100(1 - 0.6 \times 0.5) = 70$ ECU in expectation, above 60 from staying out. A monetary B would challenge. At $r = 0.9$, the value is 55 and B would stay out. The answer depends on an unobserved belief about resistance. The application does not ask for that belief. An observed B choice can therefore imply only a range of beliefs compatible with strong modeling assumptions; it does not let the instructor read the student's mind.

7.4 Exercise 3: threats, deposits, and selection

Imagine ten A players announce resistance. B challenges only two, and both resist. The calculable follow-through rate is two of two, not ten of ten; the other eight threats were not tested. Suppose also that six of the ten A players posted a deposit and both challenges targeted nondepositors. We cannot observe what depositors would have done if challenged. The deposit may have deterred, but it may also have been chosen by participants already willing to resist. Propose a contrast that assigns deposits exogenously while keeping payoffs fixed to distinguish a direct commitment effect from selection. This separates experimental treatment assignment from an observational comparison inside one treatment.

7.5 Exercise 4: Bayesian updating with noisy history

Suppose B initially regards a resistant and yielding type as equally likely. A displayed entry says "resisted" and is accurate 80% of the time. The posterior probability of a resistant type is $0.8 \times 0.5 / [0.8 \times 0.5 + 0.2 \times 0.5] = 0.8$. If the prior were 0.25, the posterior would be $0.8 \times 0.25 / [0.8 \times 0.25 + 0.2 \times 0.75]$, approximately 0.571. The record does not alter A's private cost; it alters B's information. If a student reports 0.8 for every prior, ask them to distinguish the accuracy of an observation from the posterior probability after combining it with prior odds.

7.6 Exercise 5: ex ante versus ex post

A low-cost A resists, loses the draw, and receives 25. Another A yields and receives 60. Was resistance a mistake? Not necessarily from the ex ante monetary perspective: with a 50% win probability its expected payoff was 75. Judging a decision using the realized draw introduces hindsight bias. The proper comparison uses information available on A's page before the click. If the participant believed the probability was different, the stated expectation might not match that person's assessment; the game does not measure that belief. The instructor can explain both points without attributing a motive to an individual participant.

Feedback script

"Let us first describe the choices. How many B players challenged out of how many completed encounters? Among challenged A players, how many resisted? How many had previously announced resistance? Now choose one compatible explanation and design a comparison that could test it. A single class session helps us understand the mechanism; it does not by itself establish a universal theory of international crises."

Data, privacy, and common problems

8.1 Export and statistical unit

The custom export has one row per pair and completed encounter. It includes round, treatment, matching mode, A's private cost, message, challenge, response, prior measure, actual and displayed histories where relevant, victory draw, outcome path, and payments. The two histories allow an audit that observation error did not change real action or payoff. The export supports linking encounters for internal analysis without putting personal data in the projectable report. The projectable instructor report is aggregate and does not name students. Do not publish individual links or raw records without following the course's data and privacy rules.

When estimating a repeated-game rate, say whether the denominator consists of encounters or people. One student can appear in ten rows, and their responses may be correlated. A formal comparison across sessions should respect the treatment assignment unit and repeated observations in its uncertainty calculations. The app records decisions and payments but asks no demographic or political-attitude questions. It also does not elicit B's beliefs about resistance. If research requires those measures, use a separately approved instrument and protocol rather than inferring them from actions.

8.2 Troubleshooting

If someone remains on a waiting page, first check that their partner used a different individual link and completed the preceding page. A short delay is normal in a sequential A–B game. If several links lead to the same role or state, examine whether the same row was copied repeatedly. Do not reset the database to solve a delay; that destroys active sessions and links. If a participant leaves, mark the encounter as pending in discussion and create a new, correctly sized session for a later class rather than manufacturing a response.

If aggregates are missing, verify that the instructor pressed reveal and at least one encounter is complete. A dash in a rate means a zero denominator, not a broken calculation. If a basic room offers extended treatments, stop and check the family and room filters. If a page shows duplicate large headings, ensure the instance has restarted with the current code and inform lab support; the approved design uses one principal heading per participant screen. If an individual link fails after a database reset, create a new session and distribute its new links. Do not use test-session data in a real class.

8.3 Final operational checklist

Before the session, verify server, language, room, configuration, even participant count, and distinct links. During play, maintain silence, avoid disclosing private costs, observe progress, and wait for outcomes. Before feedback, state complete and pending encounters, reveal aggregates, and read the denominators. After feedback, export only when needed, store data in an approved location, and close the instructor session. When comparing variants, write down exactly which rule changed and keep facts, interpretations, and inferential limits separate.

Comparison and inference in class

9.1 Counterfactuals and assignment

To say a measure changed B's behavior, we need to imagine how the same pair would have acted under an alternative rule. Both outcomes can never be observed for that pair at the same time. A comparison between sessions approximates this counterfactual when people are assigned comparably and instructions, timing, and classroom context do not change together with the rule. In a small class formal random assignment may be impractical, but the instructor can still state what was compared. If students play without a deposit first and with a deposit second, a difference might be caused by the device or by experience. If half the room receives one rule and half another, initial composition may matter. The teaching value increases when this uncertainty is made explicit instead of treating one observed difference as a decisive causal estimate.

Within a session, A chooses some measures. Deposits and preparation are not randomly assigned: someone who expects a challenge, has a low cost, or values keeping their word may select into them. If B challenges depositors less often, the deposit mechanism may contribute, but characteristics of the people who chose it could also explain the difference. Ask students to design a second session that separates these stories. One option would randomly require some pairs to post a deposit while holding the rest of the game fixed. Such a design would no longer study willingness to deposit; it would study the effect of being committed. Both are worthwhile questions, but they are not the same question.

9.2 What is and is not measured

The program observes actions, assigned costs, winner draws, and payoffs. It can compute challenges among completed encounters, resistance among challenged encounters, and follow-through among tested announcements. It does not observe B's confidence in a message, subjective win probabilities, A's desire for reputation, or aversion to conflict. Classroom discussion can ask about motives, but answers given after an outcome may reflect hindsight. Measuring beliefs before choices would require a separate, carefully designed elicitation task. An action alone should not be translated into a precise personal probability without strong assumptions.

Selection also affects the response data. Yielding or resisting is observed only when B challenges. If a rule deters precisely the B players who would have faced a resistant A, the challenged A players under that rule are a different subset. A lower resistance rate among challenges need not mean the rule reduced everyone's willingness to resist. It may have deterred the riskiest challenges, leaving pairs in which B expected concession. Draw the decision tree and point to the first filter, B's choice, before examining the sample of A responses. This is why denominators change across branches and why conditional rates need careful interpretation.

9.3 Welfare, conflict, and unmeasured outcomes

The game awards points to both actors but includes no harm to outsiders. The sum of A and B payoffs is a partial within-pair efficiency measure. Without challenge the sum is 160, and with concession it is also 160. Conflict produces 125 if A's cost is 5 and 105 if it is 25, regardless of who wins: the contest destroys points through the resistance and challenge costs. Yet an equal sum of 160 conceals opposite distributions, 100/60 and 60/100. "Avoid conflict" does not settle who keeps the valuable position. Nor does the point sum represent justice, legal rights, or real-world external damage. Ask whether a negotiated transfer could preserve the total while changing its distribution.

A commitment device can also have mixed effects. If a deposit prevents challenges, it may avoid a costly fight. If a challenge occurs through error or different preferences, the same device can make A resist where A would otherwise yield. Preparation may deter but costs 5 even in peace. A backing-down cost can make a threat credible without improving physical capacity, yet it reduces flexibility once a crisis unfolds. These tensions are why the instructor should inspect challenge, response, outcome path, and mean payments together, rather than celebrating a low challenge rate alone. Theory organizes possible consequences; a class displays one realization under deliberately simple parameters.

Evidence and references

10.1 How to use the literature

Crisis bargaining theory explains why hidden costs can generate challenges that common information might prevent. Research on reputation asks how past actions shape expectations, while recognizing that a past response may indicate a stable type, a particular incentive, or strategic investment in future behavior. The experimental articles below motivate mechanisms and classroom questions; they do not predict that this class will reproduce any published effect size. The automatic 20 ECU retreat penalty, 25 ECU deposit, and 20%-error history are pedagogical parameters. No single cited article supplies this exact protocol.

Tomz (2007) provides experimental evidence concerning domestic audience costs in international relations: a leader who backs down may face a political penalty. Our automatic cost represents a change in incentives but does not observe a real audience. Tingley and Walter (2011) use incentivized experiments on repeated play and reputation building, motivating the possibility that current resistance affects future challenges. Weisiger and Yarhi-Milo (2015) examine how past actions matter for reputational judgments. Their work helps explain why a history may influence choices without being a perfect measure of future resolve. In each case, distinguish the theoretical mechanism, published evidence, and the descriptive pattern from the current classroom.

Bibliography

- Tomz, Michael. 2007. "Domestic Audience Costs in International Relations: An Experimental Approach." *International Organization* 61(4): 821–840. <https://www.cambridge.org/core/journals/international-organization/article/abs/domestic-audience-costs-in-international-relations-an-experimental-approach/DE9C218365E2AF849FC6ECF67CCC45F1>.
- Tingley, Dustin, and Barbara F. Walter. 2011. "The Effect of Repeated Play on Reputation Building: An Experimental Approach." *International Organization* 65(2): 343–365. <https://www.cambridge.org/core/journals/international-organization/article/abs/effect-of-repeated-play-on-reputation-building-an-experimental-approach/0CCB307AC23930F706A8CBDC2F1BF57F>.
- Weisiger, Alex, and Keren Yarhi-Milo. 2015. "Revisiting Reputation: How Past Actions Matter in International Politics." *International Organization* 69(2): 473–495. <https://www.cambridge.org/core/journals/international-organization/article/revisiting-reputation-how-past-actions-matter-in-international-politics/8BF54EA849FC4925FDC4C43FB9A810C4>.
- American Economic Association. "JEL Classification Codes Guide." <https://www.aeaweb.org/econlit/jelCodes.php?view=econlit>.