Behavioral game theory
Appendix 2:
Extensive Form, SPNE, and Beliefs

Summer 26
Niels Boissonnet, Bielefeld

Why this appendix matters

This appendix collects the formal tools that the main lecture sequence uses repeatedly once strategic interaction is **dynamic** or **informationally rich**.

- Lectures on backward induction, centipede, signaling, promises, and auctions all rely on the extensive-form language.
- The strategic form is still there in the background, but in dynamic environments the key economic objects are often **timing**, **observability**, and **beliefs**.
- The practical question is therefore not just "what is an equilibrium?", but also **which equilibria remain plausible once players can react to what has happened before**.

Learning goals

After working through this appendix, you should be able to:

1. read and construct an extensive-form representation of a dynamic game;
2. distinguish clearly between **actions**, **information sets**, and **strategies**;
3. explain why Nash equilibrium can be too weak in dynamic environments;
4. solve finite complete-information games by **backward induction**;
5. define **subgames**, **subgame-perfect equilibrium**, and the basic idea of **perfect Bayesian equilibrium**;
6. keep separate the roles of **optimality**, **beliefs**, and **Bayesian updating**.

Route through the appendix

1. **Reading game trees:** what the extensive form adds relative to the strategic form.
2. **Complete information:** why sequential rationality leads to backward induction and subgame perfection.
3. **Incomplete information:** why strategies are not enough and why beliefs have to be added.
4. **Summary and self-check:** the main concepts you should carry into the main lecture sequence.

Interpretive thread: each refinement adds discipline because dynamic play creates opportunities for **non-credible plans**, **learning from history**, and **belief-dependent responses**.

Reading game trees

Why the strategic form is not enough on its own

The strategic form is excellent for listing players, strategy sets, and payoffs. But it suppresses several objects that become economically important in behavioral game theory:

- **Order of moves:** who acts first and who reacts later?
- **Repeated opportunities to move:** can a player act again after observing something?
- **Information flow:** what does a player know when she chooses?
- **Interpretation of equilibrium:** is a proposed continuation plan actually optimal once the relevant history has occurred?

Take-away: the extensive form makes the temporal and informational structure explicit instead of hiding it inside complete contingent strategies.

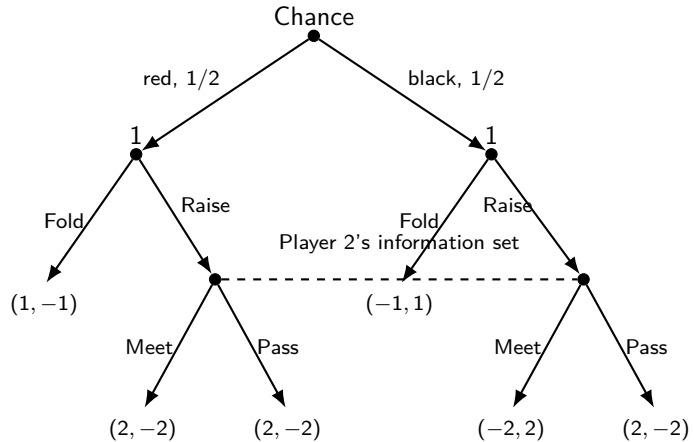
A simple card game as running example

Consider the following dynamic game.

- First, **chance** draws a card for Player 1: red with probability $1/2$, black with probability $1/2$.
- Player 1 observes the card and chooses **Raise** or **Fold**.
- If Player 1 folds, payoffs are $(1, -1)$ after a red card and $(-1, 1)$ after a black card.
- If Player 1 raises, Player 2 observes only that a raise occurred, but not the color of the card.
- Player 2 then chooses **Meet** or **Pass**. If Player 2 meets, the card is revealed and payoffs depend on the card; if Player 2 passes, Player 1 wins the pot.

Why this example is useful: it combines chance, private information, and a later response by another player.

The same game as a tree



What the tree reveals immediately: chance moves first, Player 1 knows the card, and Player 2 observes the raise but not the card color.

What an extensive-form game specifies

Definition

An **extensive-form game** specifies:

1. the set of players;
2. the **order** in which moves occur;
3. the **action set** available at each decision point;
4. the **information** available when a player moves;
5. the role of **chance**, if nature moves;
6. the **payoffs** attached to terminal histories.

Interpretation: the extensive form is not just a picture. It is a structured description of histories, opportunities to move, and information.

Nodes, branches, information sets, and histories

- A **node** is a point in the game tree. Decision nodes belong to players or to chance.
- A **branch** records one feasible move from a node.
- A **history** is a path from the root to a current node.
- A **terminal history** ends the game and determines payoffs.
- An **information set** groups decision nodes that the moving player cannot distinguish when choosing.

Economic meaning

An information set captures the idea that a player knows *that* she moves, but not *exactly where* in the tree she is.

Actions are not strategies

- An **action** is one move at one information set, for example "Raise" or "Pass".
- A **strategy** is a complete contingent plan: it tells a player what to do at *every* information set at which she could possibly move.

Example

In the card game, Player 1 moves after a red card and after a black card. So a pure strategy for Player 1 must specify one action after red and one action after black, for example

(Raise after red, Fold after black).

Why this matters: equilibrium concepts are defined over **strategies**, but sequential rationality will later evaluate whether the implied **actions** are sensible after a given history.

Perfect and imperfect information

Perfect information

Every information set is a singleton.

- each player observes all earlier moves;
- there is no private hidden action by another player;
- backward induction can be applied directly in finite games.

Imperfect information

Some information set contains several nodes.

- some previous move or state is unobserved;
- later play requires **beliefs** about where the game is;
- subgame perfection is usually not enough.

Caution: "imperfect information" and "incomplete information" are not synonyms. Imperfect information is about what is observed during play; incomplete information is about uncertainty over payoff-relevant features such as types.

Extensive form and strategic form

Every finite extensive-form game can be translated into a strategic-form game:

- list every player's complete contingent plans;
- treat those complete plans as the players' pure strategies;
- compute payoffs induced by the terminal history generated by the strategy profile.

Why we still keep the extensive form

The translation preserves strategic possibilities, but it hides the dynamic logic:

- which continuation plans are incredible,
- which histories create learning,
- which information sets require beliefs,
- and which solution concept is economically appropriate.

Take-away from Part I

- The extensive form is the right language once timing and observability matter.
- The key new objects are **histories** and **information sets**.
- An action is a local move; a strategy is a full contingent plan.
- The strategic form remains available, but the extensive form is needed to evaluate dynamic plausibility.

Complete Information and Subgame Perfection

Why Nash equilibrium can be too weak

Consider a simple entry-deterrence game:

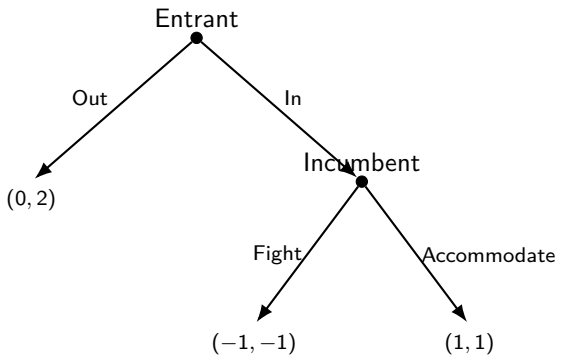
- Entrant chooses **In** or **Out**.
- If Entrant chooses **Out**, payoffs are $(0, 2)$.
- If Entrant chooses **In**, Incumbent chooses **Fight** or **Accommodate**.
- After entry, payoffs are $(-1, -1)$ under Fight and $(1, 1)$ under Accommodate.

	Fight	**Accommodate**
In	$-1,\ -1$	$1,\ 1$
Out	$0,\ \ 2$	$0,\ \ 2$

In the strategic form, "Out" by the entrant and "Fight if entry occurs" by the incumbent is a Nash equilibrium: the entrant stays out because of the threat to fight.

Problem: once entry actually occurs, **Fight** is no longer optimal for the incumbent.

Entry deterrence as a tree



Economic lesson: the extensive form lets us ask whether continuation behavior is optimal *at the point where it must actually be chosen*.

Sequential rationality

Core requirement

A strategy profile is **sequentially rational** if each player's prescribed action is optimal whenever the corresponding information set is reached.

In complete-information games with singleton information sets, this means:

- after every possible history, the continuation plan must still make sense;
- a player cannot be asked to carry out a continuation action that is obviously against her own interest when the time comes;
- non-credible threats are excluded.

Thus: dynamic reasoning imposes stronger discipline than Nash equilibrium in strategic form.

Subgames and subgame perfection

Definition

A **subgame** is a part of the game tree that starts at a decision node that is itself a singleton information set and that contains all successor nodes of that decision node.

Definition

A strategy profile is a **subgame-perfect equilibrium** (SPNE) if it induces a Nash equilibrium in the whole game and in every subgame.

Interpretation: SPNE requires equilibrium behavior not only from the beginning of the game, but also after every history that can start a well-defined continuation game.

Backward induction in finite games

In finite perfect-information games, SPNE can be found by **backward induction**:

1. Look at the final decision nodes.
2. At each such node, determine the moving player's optimal action.
3. Replace the final subgames by their equilibrium payoffs.
4. Move one step backward and repeat.
5. Continue until the initial node is reached.

What backward induction uses

Future continuation play is treated as rational, so current players choose the best response to those rational continuations.

Backward induction in the entry game

Apply backward induction to the tree:

- If entry occurs, the incumbent compares **Fight** yielding -1 to **Accommodate** yielding 1.
- So the incumbent chooses **Accommodate**.
- Anticipating this, the entrant compares **Out** yielding 0 to **In** yielding 1.
- Hence the entrant chooses **In**.

Conclusion

The unique SPNE is **In** by the entrant and **Accommodate** by the incumbent.

Lesson: the equilibrium that relied on the threat to fight is Nash, but not subgame perfect.

Application: Stackelberg quantity competition

Two firms choose quantities sequentially.

- Firm 1 (leader) chooses $q_1 \geq 0$ first.
- Firm 2 (follower) observes q_1 and then chooses $q_2 \geq 0$.
- Inverse demand is $P(Q) = a - Q$ for $Q = q_1 + q_2$.
- Both firms have constant marginal cost $c < a$.
- Profit of firm i is $\pi_i = (P(Q) - c)q_i$ whenever $P(Q) \geq c$.

Question: how does the leader exploit the follower's optimal reaction?

Solving Stackelberg by backward induction

Step 1: solve the follower's problem for given q_1:

$$\pi_2(q_1, q_2) = (a - q_1 - q_2 - c)q_2.$$

The best reply satisfies

$$R_2(q_1) = \frac{a - c - q_1}{2}.$$

Step 2: the leader anticipates that reaction and chooses

$$q_1^* \in \arg\max_{q_1 \geq 0} \ (a - q_1 - R_2(q_1) - c)q_1 = \arg\max_{q_1 \geq 0} \frac{a - c - q_1}{2} q_1.$$

Hence

$$q_1^* = \frac{a - c}{2}, \qquad q_2^* = R_2(q_1^*) = \frac{a - c}{4}.$$

Economic interpretation: the leader commits early and benefits from the follower's reaction.

What SPNE adds in complete-information games

- It rules out equilibria supported only by non-credible continuation plans.
- It fits naturally with environments in which later players observe earlier moves.
- It produces a conceptually clean benchmark for many behavioral tests:
  - backward induction in centipede,
  - deterrence and reputation,
  - sequential bargaining and market entry,
  - dynamic oligopoly.
- In finite perfect-information games, backward induction is the practical solution method.

Incomplete Information and Perfect Bayesian Equilibrium

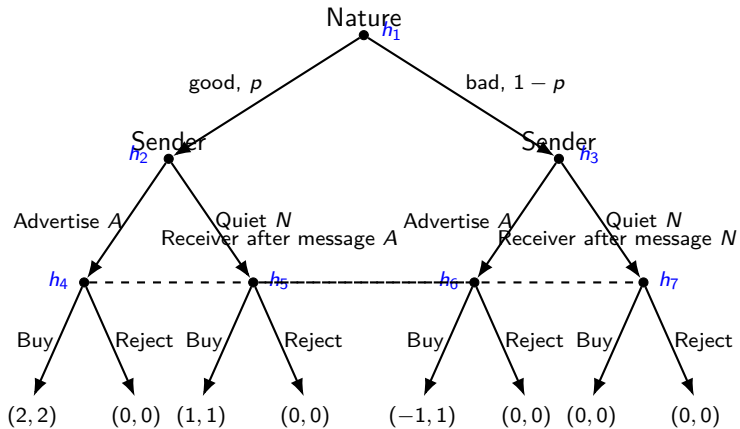
Why SPNE is not enough once information is imperfect

Subgame perfection works well when continuation games start at singleton information sets. It becomes incomplete once a player moves at an information set containing several nodes.

- The moving player must decide without knowing the exact history.
- To choose optimally, she needs a **belief** over the nodes in the information set.
- The equilibrium concept must therefore specify both **strategies** and **beliefs**.

Core idea: with hidden information, rational continuation behavior is impossible to evaluate without also specifying what players think has happened.

A simple signaling game



Nature draws the sender's type. The sender observes the type and sends a message. The receiver observes the message but not the type, and then chooses whether to buy.

Mixed strategies in extensive-form games

In extensive-form games, a mixed or behavioral strategy assigns a probability distribution over available actions at each information set.

- In finite games with perfect recall, mixed and behavioral strategies are outcome-equivalent.
- For equilibrium work in dynamic games, it is often most convenient to write **behavioral strategies**: at each information set, directly specify the probabilities of the feasible actions there.
- Beliefs are then updated from the probabilities induced by these local action choices.

Why this is useful: many PBE examples involve on-path randomization and off-path inference.

Mixed Strategies

Let $I : N \mapsto 2^N$ a function that associates each node with its information set, e.g.

- $I(h_2) = \{h_2\}$ no uncertainty
- $I(h_4) = \{h_4, h_6\}$ uncertainty

A mixed strategy σ_i assigns to each information set a prob over action.

- $\sigma_1^{I(h_2)}(A)$ probability that P1 chooses A given that she knows that she is in h_2
- $\sigma_2^{I(h_4)}(B)$ probability that P2 chooses *buy* given that she believes she is either in h_4 or h_6

Beliefs at information sets

Suppose a player's information set I contains nodes $h_1, \ldots, h_k$. A **belief** at I is a probability distribution

$$\mu(\cdot \mid I) \in \Delta(\{h_1, \ldots, h_k\}),$$

where $\mu(h_j \mid I)$ is the player's subjective probability that the true history is h_j.

Interpretation

Beliefs summarize how the player interprets the information revealed by earlier play.

Beliefs and Bayesian updating

The Prior belief μ_0 = the Players' beliefs before the game start.

- P2 believes nature has drawn "good" with prob $\mu_0(g) = p$
- Usually, common prior is assumed, i.e. both players have the same prior

Beliefs and Bayesian updating

The Prior belief μ_0 = the Players' beliefs before the game start.

- P2 believes nature has drawn "good" with prob $\mu_0(g) = p$
- Usually, common prior is assumed, i.e. both players have the same prior

The Posterior beliefs = the Players' beliefs after observing some moves

- After observing P1 advertized (at $I = \{h_4, h_6\}$), P2 updates her belief $\mu_2(h_4|I)$

$$\mu_2(h_4|I) = \frac{\mathbb{P}(I|h_4)\mu_0(h_4)}{\mathbb{P}(I|h_4)\mu_0(h_4) + \mathbb{P}(I|b)\mu_0(b)}$$

Beliefs and Bayesian updating

The Prior belief μ_0 = the Players' beliefs before the game start.

- P2 believes nature has drawn "good" with prob $\mu_0(g) = p$
- Usually, common prior is assumed, i.e. both players have the same prior

The Posterior beliefs = the Players' beliefs after observing some moves

- After observing P1 advertized (at $I = \{h_4, h_6\}$), P2 updates her belief $\mu_2(h_4|I)$

$$\mu_2(h_4|I) = \frac{\mathbb{P}(I|h_4)\mu_0(h_4)}{\mathbb{P}(I|h_4)\mu_0(h_4) + \mathbb{P}(I|b)\mu_0(b)}$$

Beliefs Consistency= the bielefs must be consistent with the strategies played

$$\mu_2(h_4|I) = \frac{\sigma_1^{\{h_2\}}(A)\mu_0(h_2)}{\sigma_1^{\{h_2\}}(A)\mu_0(h_2) + \sigma_1^{\{h_3\}}(A)\mu_0(h_3)}$$

Bayesian updating on the path of play

A strategy profile induces probabilities on histories. Whenever an information set is reached with positive probability, beliefs should be obtained from those induced probabilities by **Bayes' rule**.

$$\mu(h \mid I) = \frac{\Pr_\sigma(h)}{\sum_{\tilde{h} \in I} \Pr_\sigma(\tilde{h})} \qquad \text{for each } h \in I \text{ with } \sum_{\tilde{h} \in I} \Pr_\sigma(\tilde{h}) > 0.$$

On-path idea: beliefs should be consistent with the equilibrium strategy profile wherever that strategy profile actually leads.

Off-path issue: if an information set is reached with probability zero, Bayes' rule does not pin down beliefs uniquely. Different refinements then add more discipline.

Perfect Bayesian equilibrium

Definition

A **perfect Bayesian equilibrium** (PBE) consists of

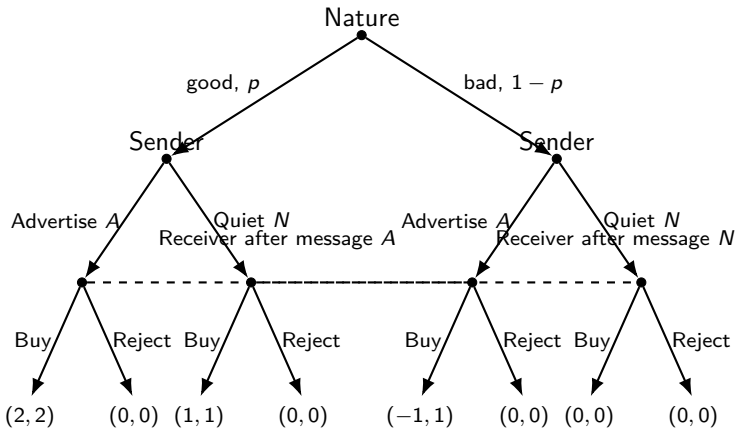
- a strategy profile σ, and
- a system of beliefs μ at every information set,

such that

1. each player's strategy is sequentially rational given the beliefs and the other players' strategies;
2. beliefs are consistent with Bayes' rule at every information set reached with positive probability.

Interpretation: PBE combines dynamic optimality with disciplined inference from observed play.

A simple signaling game



Nature draws the sender's type. The sender observes the type and sends a message. The receiver observes the message but not the type, and then chooses whether to buy.

How to reason in the signaling game

To test whether a candidate equilibrium is a PBE, ask three separate questions.

1. **Sender incentives:** given how the receiver reacts to each message, which message does each sender type prefer?
2. **Receiver beliefs:** after observing a message, what posterior probability does the receiver assign to the sender being good?
3. **Receiver optimality:** given those beliefs, is buying or rejecting optimal after each message?

Typical structures:

- **Pooling:** both types choose the same message, so the receiver learns little.
- **Separating:** different types choose different messages, so the message reveals the type.
- **Semi-separating / mixed:** types randomize or separate only partially.

Where beliefs come from in signaling games

Suppose both types choose message A in equilibrium. Then after observing A, the receiver's posterior equals the prior:

$$\mu(\text{good} \mid A) = p.$$

If instead only the good type chooses A, then Bayes' rule implies

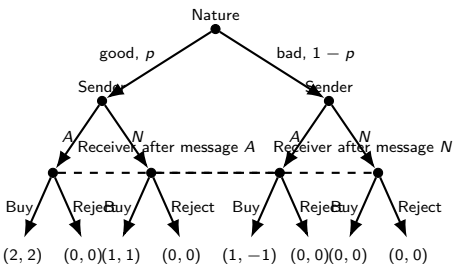
$$\mu(\text{good} \mid A) = 1.$$

Why this matters

The receiver's action after a message depends on what the message reveals about type. Thus the sender's strategic problem and the receiver's inference problem are tightly linked.

Main insight: dynamic games with incomplete information are solved jointly in terms of **actions** and **interpretations of observed actions**.

Receiver's best response



Given belief μ_A after message A:

$$EU_R(\text{Buy} \mid A) = 2\mu_A - (1 - \mu_A) = 3\mu_A - 1$$

$$\text{Buy optimal} \iff \mu_A \geq \tfrac{1}{3}$$

Given belief μ_N after message N:

$$EU_R(\text{Buy} \mid N) = \mu_N$$

$$\text{Buy optimal} \iff \mu_N > 0 \quad \text{(indifferent if } \mu_N = 0\text{)}$$

Equilibrium 1: Pooling on N

Strategies: both types send N.

Beliefs: $\mu_N = p$ (on path); $\mu_A = 0$ (off path).

Receiver: Reject after A (since $\mu_A = 0 < 1/3$); Buy after N (since $p > 0$).

Checking incentives:

- *bad*: N gives 0; deviating to A gives 0 (Rejected) $\Rightarrow$ no gain.
- *good*: N gives 1; deviating to A gives 0 (Rejected) $\Rightarrow$ no gain.

Result

Valid PBE for *any* $p \in (0, 1)$.

Equilibrium 2: Separating

Strategies: good $\to A$, bad $\to N$.

Beliefs: $\mu_A = 1$, $\mu_N = 0$ (both pinned down by Bayes' rule).

Receiver: Buy after A (since $\mu_A = 1 \geq 1/3$); Buy after N (selecting the indifferent action).

Checking incentives:

- *good*: A gives 2; deviating to N gives 1 $\Rightarrow$ stays at A.
- *bad*: N gives 0; deviating to A gives -1 $\Rightarrow$ stays at N.

Result

Valid PBE for *any* $p \in (0, 1)$.

Another example

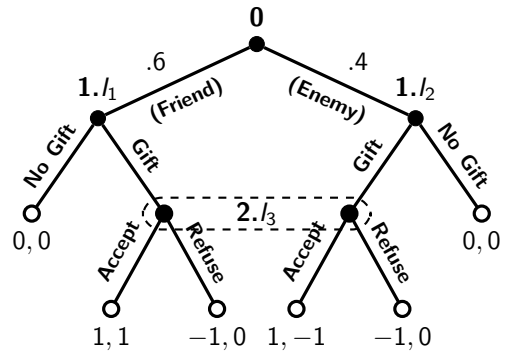
P1 is a friend of P2 (prob 0.6) or an enemy;

He is the only one to know it and must first decide to offer a gift to P2 or not;

If no gift then both get 0;

If yes P2 can accept or refuse the gift

- **Refusal:** P1 gets -1 and P2 0;
- **Acceptance:** P1 gets 1 and P2 gets 1 if Player 1 is a friend, and she gets -1 otherwise.



Strategies

- **Player 1:** GG, GN, NG, NN
- **Player 2:** A, R

Obtaining Perfect-Bayesian Equilibrium

$$EU_2(\sigma_2^{I_3}) = (\mu_f - \mu_e)\sigma_2^{I_3}(a)$$
$$EU_1(\sigma_1^{I_i}) = \sigma_1^{I_i}(g)(2\sigma_2^{I_3}(a) - 1) \qquad \text{for } i = 1, 2$$

The three following equilibria are possible,

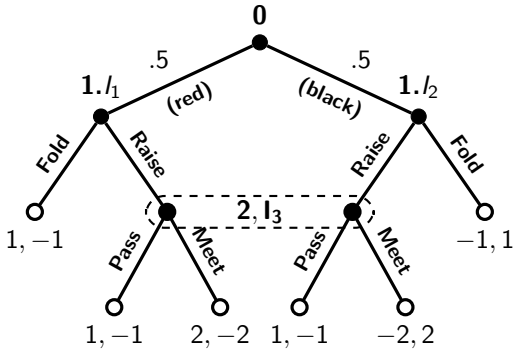
Obtaining Perfect-Bayesian Equilibrium

$$EU_2(\sigma_2^{I_3}) = (\mu_f - \mu_e)\sigma_2^{I_3}(a)$$
$$EU_1(\sigma_1^{I_i}) = \sigma_1^{I_i}(g)(2\sigma_2^{I_3}(a) - 1) \qquad \text{for } i = 1, 2$$

The three following equilibria are possible,

1. $\sigma_1^{I_i}(g) = 0$ for both $i = 1, 2$, I_3 is off-path, so μ_f is unrestricted by Bayes' rule. Any $\mu_f < 0.4$ gives $\sigma_2^{I_3}(a) = 0$, which makes "no gift" optimal for both types — consistent.
2. if $\sigma_2^{I_3}(a) \neq 1/2$, then $\sigma_1^{I_1}(g) = \sigma_1^{I_2}(g) \in \{1, 0\}$ (pooling). If both gift ($= 1$):
$$\mu_f = \frac{1 \times 0.6}{1 \times 0.6 + 1 \times 0.4} = 0.6 \qquad \text{implying } \sigma_2^{I_3}(a) = 1 \neq 1/2 \text{ — consistent.}$$
3. if $\sigma_2^{I_3}(a) = 1/2$, then it must be that
$$\mu_f = \mu_e \implies \frac{\sigma_1^{I_1}(g)\,0.6}{\sigma_1^{I_1}(g)\,0.6 + \sigma_1^{I_2}(g)\,0.4} = \frac{\sigma_1^{I_2}(g)\,0.4}{\sigma_1^{I_1}(g)\,0.6 + \sigma_1^{I_2}(g)\,0.4} \implies \sigma_1^{I_2}(g) = \frac{3}{2}\sigma_1^{I_1}(g)$$
for any $\sigma_1^{I_1}(g) \in (0, 2/3]$.

A last example



- Player 1's information sets I_1="the card is red", or I_2="the card is black";
- Player 1's information sets I_3="Player 1 has chosen Raise";

Let's find the PBE

P2's expeected utility: We have $EU_2(\sigma_2^{l_3}) = \sigma_2^{l_3}(m) - 1 + \sigma_2^{l_3}(m)(2\mu_b - 2\mu_r)$, so

Let's find the PBE

P2's expeected utility: We have $EU_2(\sigma_2^{I_3}) = \sigma_2^{I_3}(m) - 1 + \sigma_2^{I_3}(m)(2\mu_b - 2\mu_r)$, so

$$EU_2(\sigma_2^{I_3}) = \sigma_2^{I_3}(m)\phi(\mu) - 1 \qquad \text{where } \phi(\mu) = 1 + 2\mu_b - 2\mu_r$$

$$\text{And, the maximizing } \sigma_2^{I_3}(m) = \begin{cases} 0 & \text{if } \phi(\mu) > 0 \\ \text{anything} & \text{if } \phi(\mu) = 0 \\ 1 & \text{if } \phi(\mu) > 0 \end{cases}$$

P1's expected utility

$$EU_1(\sigma_1^{I_1}) = \sigma_1^{I_1}(r)(2\sigma_2^{I_3}(m) + 1 - \sigma_2^{I_3}(m)) + 1 - \sigma_1^{I_1}(r) = \sigma_1^{I_1}(r)\sigma_2^{I_3}(m) + 1$$
$$EU_1(\sigma_1^{I_2}) = \sigma_1^{I_2}(r)(-2\sigma_2^{I_3}(m) + 1 - \sigma_2^{I_3}(m)) - 1 + \sigma_1^{I_2}(r) = \sigma_1^{I_2}(r)(2 - 3\sigma_2^{I_3}(m)) + 1$$

Let's find the PBE

P2's expeected utility: We have $EU_2(\sigma_2^{I_3}) = \sigma_2^{I_3}(m) - 1 + \sigma_2^{I_3}(m)(2\mu_b - 2\mu_r)$, so

$$EU_2(\sigma_2^{I_3}) = \sigma_2^{I_3}(m)\phi(\mu) - 1 \qquad \text{where } \phi(\mu) = 1 + 2\mu_b - 2\mu_r$$

$$\text{And, the maximizing } \sigma_2^{I_3}(m) = \begin{cases} 0 & \text{if } \phi(\mu) > 0 \\ \text{anything} & \text{if } \phi(\mu) = 0 \\ 1 & \text{if } \phi(\mu) > 0 \end{cases}$$

P1's expected utility

$$EU_1(\sigma_1^{I_1}) = \sigma_1^{I_1}(r)(2\sigma_2^{I_3}(m) + 1 - \sigma_2^{I_3}(m)) + 1 - \sigma_1^{I_1}(r) = \sigma_1^{I_1}(r)\sigma_2^{I_3}(m) + 1$$
$$EU_1(\sigma_1^{I_2}) = \sigma_1^{I_2}(r)(-2\sigma_2^{I_3}(m) + 1 - \sigma_2^{I_3}(m)) - 1 + \sigma_1^{I_2}(r) = \sigma_1^{I_2}(r)(2 - 3\sigma_2^{I_3}(m)) + 1$$

First remark: in equilibrium it can't be that $\sigma_1^{I_1}(r) = \sigma_1^{I_2}(r) = 0$ Indeed, $\sigma_1^{I_1}(r) = 0$ optimal only if $\sigma_1^{I_2}(r) = 0$ which implies that $\sigma_1^{I_2}(r) = 1$ is optimal.

Non-equilibria

- If $\sigma_2^{I_3}(m) = 1$, then it is optimal that $\sigma_1^{I_1}(r) = 1$ and $\sigma_1^{I_1}(r) = 0$. Thus,

$$\mu_r = 1 \qquad \mu_b = 0 \text{ so that } \phi(\mu) > 0 \qquad \textbf{Contradiction}$$

- If $\sigma_2^{I_3}(m) = 0$ (i.e. if $\phi(\mu) > 0$, then any $\sigma_1^{I_2}(r) = 1$. Hence,

$$1 + 2\frac{\sigma_1^{I_2}(r)\frac{1}{2}}{\sigma_1^{I_2}(r)\frac{1}{2} + \sigma_1^{I_1}(r)\frac{1}{2}} \leq 2\frac{\sigma_1^{I_1}(r)\frac{1}{2}}{\sigma_1^{I_2}(r)\frac{1}{2} + \sigma_1^{I_1}(r)\frac{1}{2}}$$

$$\implies 1 = \sigma_1^{I_2}(r) \leq \frac{1}{3}\sigma_1^{I_1}(r) \implies 3 \leq \sigma_1^{I_1}(r) \qquad \textbf{Contradiction}$$

The unique equilibrium

The only remaining possibility therefore implies $\phi(\mu) = 0$, i.e. $\sigma_1^{I_2}(r) = \frac{1}{3}\sigma_1^{I_1}(r)$.

Claim: At the equilibrium, $\phi(\mu) = 0$ entails $0 < \sigma_1^{I_2}(r) < \sigma_1^{I_1}(r) = 1$.

- By the previous remark, in equilibrium, we can't have $\sigma_1^{I_1}(r) = \sigma_1^{I_2}(r) = 0$.
- Thus both $\sigma_1^{I_1}(r)$ and $\sigma_1^{I_2}(r)$ are striclty positive since $\sigma_1^{I_2}(r) = \frac{1}{3}\sigma_1^{I_1}(r)$
- Moreover, $\sigma_1^{I_2}(r) = \frac{1}{3}\sigma_1^{I_1}(r)$, thus implies $\sigma_1^{I_2}(r) < \sigma_1^{I_1}(r)$.
- Notice that to be optimal for P1 we must have

$$\text{either} \qquad \sigma_1^{I_2}(r) \in \{0, 1\} \qquad \text{or } \sigma_1^{I_1}(r) = 1$$

Hence, the only possibility is $\sigma_1^{I_1}(r) = 1$ □

Conclusion: The claim implies $\sigma_2^{I_3}(m) = \frac{2}{3}$—otherwise $\sigma_1^{I_1}(r) = 0$ or $\sigma_1^{I_1}(r) = 1 = \sigma_1^{I_2}(r)$ and

$$\mu_r = \frac{3}{4}.$$

Computing the Subgame perfect bayesian equilibrium

Solution The SGPE is (σ_1^*, σ_2^*) with beliefs $\mu = (\mu_1, \mu_2)$ where

- $\sigma_1^{I_1}(r) = 1$ and $\sigma_1^{I_3}(r) = 1/3 \qquad \sigma_2^{I_3}(m) = 2/3$
- $\mu_2(r|I_3) = 3/4$

μ is compatible with Bayes' rule

$$\mu_2(r|I_3) = \frac{1/2\sigma_1^{I_1}(r)}{1/2\sigma_1^{I_1}(r) + 1/2\sigma_1^{I_2}(r)} = \frac{3}{4}$$

σ optimal actions in all I

- $EU_1(\sigma_1|I_1) = 2(1 + 2\sigma_1^{I_1}(r))$, max= $\sigma_1^{I_1}(r) = 1$

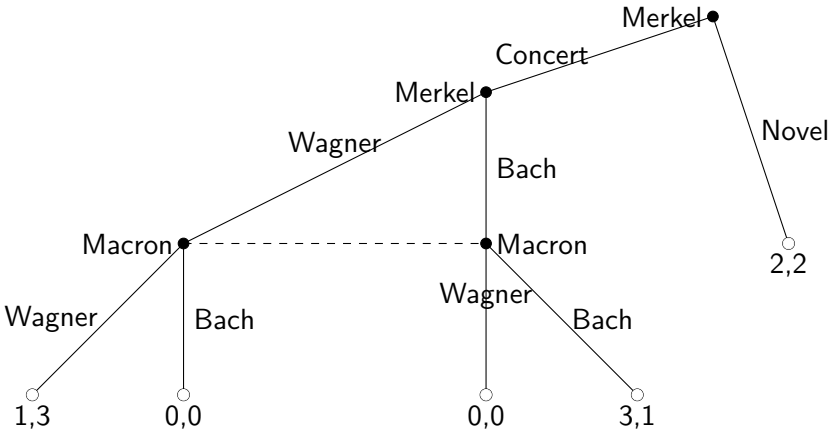
  $$\mu_2(r|I_3) = \frac{1}{1 + \sigma_1^{I_2}(r)} \implies \sigma_1^{I_2}(r) = \frac{1 - \mu_2(r|I_3)}{\mu_2(r|I_3)}$$

- $EU_1(\sigma_1|I_2) = \sigma_1^{I_2}(r)(2 - 3\sigma^{I_3}(m)) - 1 = -1$
  $\implies EU_1(\sigma_1|I_2) \geq EU_1(\sigma|I_2)$ for all σ
- $EU_2(\sigma_2|I_3) = \sigma_2^{I_3}(m)(3 - 4\mu_2(r|I_3)) - 1$
  $\implies EU_2(\sigma_3|I_3) \geq \pi(\sigma|I_3) = -1$ for all σ

Backward Induction vs Forward Induction

- **Backward Induction:** Agents believe that all future plays will be rational.
- **Forward Induction:** Agents believe that all past plays were rational.

Example:


Merkel
Concert
Merkel
Novel
Wagner
Bach
Macron
Macron
2,2
Wagner
Bach
Wagner
Bach
1,3
0,0
0,0
3,1

Problem

Subgame 1:

	W	B
CW	$3,1$	$0,0$
CB	$0,0$	$1,3$
NW	$2,2$	$2,2$
NB	$2,2$	$2,2$

Subgame 2:

	W	B
CW	$3,1$	$0,0$
CB	$0,0$	$1,3$

The Subgame-Perfect Equilibria are (CW, W) and (NB, B). Are they both reasonable?

- Imagine that Macron finds himself playing to the second subgame. Is it reasonable to believe that Merkel has chosen Bach? Probably not. Merkel could have guaranteed herself 2 by playing N. Why would she enter a subgame and play in a manner in which she is only going to get 1?
- Arguably (NB, B) is not reasonable, despite being a subgame Perfect-equilibrium.

This is an example of forward induction reasoning.

What this leaves for later lectures

We have exposed the basic toolkit, but several stronger concepts appear later:

- **Sequential equilibrium:** a sharper consistency requirement for beliefs and strategies.
- **Psychological and belief-dependent games:** where utility depends directly on beliefs or intentions.

So: PBE is the right baseline once beliefs matter, but it is not the end of the refinement story.

Take-away

- Once a player moves at a non-singleton information set, beliefs become indispensable.
- PBE combines **sequential rationality** with **Bayesian consistency on path**.
- Signaling environments are the canonical application: messages influence actions because they influence beliefs.
- Off-path beliefs matter, which is why stronger refinements can change conclusions.

Summary: what to remember

1. The extensive form is the natural representation of dynamic games.
2. The strategic form can encode the same possibilities, but it hides timing and information.
3. SPNE strengthens Nash by imposing credible continuation behavior in every subgame.
4. Backward induction is the practical method in finite perfect-information games.
5. PBE extends sequential reasoning to environments with hidden information by adding beliefs.
6. In behavioral game theory, these concepts are benchmarks against which observed play is compared.

How this appendix supports the course

- In **centipede** and **backward-induction** experiments, the key benchmark is SPNE via backward induction.
- In **cascades**, **signaling**, and **promises**, the key benchmark adds beliefs and Bayesian updating.
- In **auctions**, the extensive form clarifies which information bidders have and when.
- In **behavioral applications**, observed deviations are only meaningful once the benchmark dynamic structure is explicit.

Bottom line: learning the extensive-form language is not an isolated technical exercise; it is what makes later empirical and behavioral comparisons interpretable.

References

- Fudenberg, D. and J. Tirole (1991). *Game Theory*. MIT Press.
- Kreps, D. M. and R. Wilson (1982). Sequential equilibria. *Econometrica* 50(4), 863–894.
- Mas-Colell, A., M. D. Whinston, and J. R. Green (1995). *Microeconomic Theory*. Oxford University Press.
- Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*. MIT Press.
- Rasmusen, E. (2007). *Games and Information*. Blackwell.

Repetition Questions

Repetition questions I

1. **Why do we need the extensive form if we already have strategic form?**
   Answer sketch: because dynamic environments require explicit timing, observability, and beliefs; the strategic form hides these inside contingent plans.

2. **What is an information set?**
   Answer sketch: a set of decision nodes that the moving player cannot distinguish when choosing.

3. **What is the difference between an action and a strategy?**
   Answer sketch: an action is a move at one information set; a strategy is a complete contingent plan specifying actions at all information sets.

4. **What is the economic problem with non-credible threats?**
   Answer sketch: they can support Nash equilibria in strategic form even though the threatened action would not actually be optimal once the relevant history occurs.

Repetition questions II

5. **How does backward induction work?**
   Answer sketch: solve the final decision problems first, replace them by equilibrium payoffs, and continue recursively backward to the initial node.

6. **What is a subgame-perfect equilibrium?**
   Answer sketch: a strategy profile that is a Nash equilibrium in the whole game and in every subgame.

7. **Why is SPNE often not enough in games with imperfect information?**
   Answer sketch: because players at non-singleton information sets need beliefs about where they are, and SPNE by itself does not specify those beliefs.

8. **What are the two core ingredients of PBE?**
   Answer sketch: sequential rationality of strategies and beliefs that follow Bayes' rule whenever the information set is reached with positive probability.

Repetition questions III

9. **What is the role of Bayes' rule in dynamic games?**
   Answer sketch: it updates beliefs from equilibrium-induced probabilities on histories at information sets reached on the equilibrium path.

10. **What is the difference between pooling and separating in signaling games?**
    Answer sketch: in a pooling equilibrium both types choose the same message, whereas in a separating equilibrium different types choose different messages and the receiver can infer type from the message.

11. **Why are off-path beliefs important?**
    Answer sketch: because deviations may lead to information sets that are not reached in equilibrium, and conclusions about continuation behavior depend on how those deviations are interpreted.

12. **How does this appendix connect to behavioral game theory?**
    Answer sketch: SPNE and PBE provide the benchmark predictions against which behavioral deviations in dynamic and informational games are assessed.