IOSR Journal Of Economics and Finance (IOSR-JEF)
e-ISSN: 2321-5933, p-ISSN: 2321-5925.Volume 17, Issue 3 Ser. 2 (May. – June. 2026), PP 63-78
www.Iosrjournals.Org

Quantitative Economics: Game Theory And The Micro Foundations Of Strategic Behaviour

Author

Abstract
This paper introduces game theory — the mathematical study of strategic decision-making — as the modern micro foundation of economic analysis. Beginning from the assumptions of rational choice and the construction of utility functions, the paper develops the central concept of Nash equilibrium in both pure and mixed strategies, and then extends the framework to dynamic games, repeated interactions, and games of incomplete information. Each concept is illustrated with a worked quantitative example: the Cournot duopoly model is solved explicitly to derive equilibrium quantities and prices; a simulated tournament of the Iterated Prisoner's Dilemma demonstrates the surprising success of conditional cooperation; and Akerlof's lemons model is used to show how asymmetric information can cause markets to unravel. The paper closes with a survey of experimental evidence — most notably from ultimatum-game studies across cultures — that complicates the standard rational-agent assumption and motivates the rise of behavioural economics. Throughout, the goal is to show that the abstract apparatus of game theory is not a mathematical curiosity but a working language for thinking about markets, contracts, politics, and everyday life.

Date of Submission: 16-05-2026 Date of Acceptance: 26-05-2026

I. Introduction

Why economics needs game theory

For most of its history, economics was a discipline of *non-strategic* decision-making. A consumer chose a basket of groceries given fixed prices; a firm chose how much wheat to grow given the going market rate. The mathematical apparatus was that of constrained optimisation — calculus applied to utility functions and production functions — but the agents in these models did not have to think about each other. Prices were taken as given, parametric, anonymous.

The trouble with this picture is that most of the economic decisions that really matter are not like that at all. When a pharmaceutical company decides whether to launch a new drug, it must anticipate how its rival will price its own competing product. When a country decides whether to impose tariffs, it must anticipate retaliation. When a buyer haggles over a used car, the price that emerges depends on what each party believes the other knows and is willing to do. The interesting economic problems are *strategic*: each agent's best move depends on what the others are likely to do, and what the others are likely to do depends, in turn, on what they think we will do.

Game theory is the branch of mathematics — and now the branch of economics — that takes this circularity seriously. Pioneered by John von Neumann and Oskar Morgenstern in their 1944 treatise *Theory of Games and Economic Behavior*, and given its central solution concept by John Nash in 1950, game theory provides a precise language for thinking about strategic interaction. In the seventy-five years since Nash's paper, it has become so fundamental that it is now widely described as the **microfoundation** of modern economics — the layer of theory upon which everything from competition policy to mechanism design to behavioural economics is built.

What this paper does

This paper has three goals. The first is expository: to explain the central tools of game theory in language and mathematics accessible to a high-school reader who has met algebra, basic probability, and the idea of a function. The second is quantitative: every major concept is illustrated with a worked numerical example, often accompanied by a figure showing how the model behaves. The third is critical: the paper closes by surveying experimental and empirical evidence which suggests that real human beings do not always play the equilibria the theory predicts, and asks what we should make of that.

The structure is roughly as follows. Chapter 2 sets out the microfoundations — rational choice, utility, and expected utility — that game theory rests on. Chapters 3 to 6 develop the theory of static games of complete information: payoff matrices, dominant strategies, Nash equilibrium, mixed strategies, and the Cournot model of oligopoly. Chapters 7 and 8 introduce dynamic games and repeated games, with a simulated Axelrod-style tournament showing how cooperation can emerge. Chapters 9 and 10 cover games with asymmetric information — adverse selection, signalling — and auction theory. Chapter 11 turns to the experimental evidence, with a focus on the ultimatum game. A short concluding chapter draws the threads together.

A note on the mathematics

The mathematics required is modest. The reader will need to be comfortable with algebraic manipulation, basic probability (expected values), and the idea that a derivative measures a rate of change. No multivariable calculus, no measure theory, no fixed-point theorems are used in the main exposition; where they would normally appear (most notably in the proof that every finite game has at least one Nash equilibrium), the result is simply stated and illustrated. The aim is not to train a professional game theorist in twenty-five pages — an impossible task — but to convey the structure of the discipline and to communicate why economists find it so powerful.

II. Microfoundations: Rational Choice And Utility

Before we can ask how people behave when their decisions interact, we need a theory of how a single person makes decisions in isolation. That theory is **rational choice**, and it is the foundation on which all of game theory is built. The basic claim is disarmingly simple: an agent has preferences over possible outcomes, those preferences are consistent in a particular technical sense, and the agent chooses whichever feasible outcome ranks highest.

Preferences and utility

Let X denote the set of possible outcomes — a set of consumption bundles, lottery tickets, career choices, vacation destinations, whatever the agent must choose among. A **preference relation** $\succcurlyeq$ over X is a binary relation: writing $x \succcurlyeq y$ means "the agent finds x at least as good as y". Rational choice theory imposes two axioms:

- Completeness. For any pair x, y in X, either $x \succcurlyeq y$ or $y \succcurlyeq x$ (or both, in which case the agent is indifferent). The agent has an opinion, even if a weak one, about every pair.
- Transitivity. If $x \succcurlyeq y$ and $y \succcurlyeq z$, then $x \succcurlyeq z$. Preferences do not loop.

These two axioms are far less innocent than they look — completeness in particular has been contested by philosophers and behavioural economists — but together they imply a beautiful representation theorem: there exists a real-valued function $u : X \to \mathbb{R}$, called a **utility function**, such that $x \succcurlyeq y$ if and only if $u(x) \geq u(y)$. The function u is not unique; any strictly increasing transformation of u represents the same preferences. Utility, in the ordinal sense, is not a quantity of "happiness" in some absolute scale. It is just a numerical re-labelling of the agent's preference ordering, useful because it lets us do calculus and optimisation.

Definition. An agent is rational in the sense of standard microeconomics if her preferences over outcomes satisfy completeness and transitivity, and if she chooses, from her feasible set, the option(s) that maximise her utility.

Optimisation under constraint: a worked example

Consider a consumer with two goods, x and y, and the Cobb–Douglas utility function $u(x, y) = x^{½} y^{½}$. Let her budget constraint be $2x + y = 8$ (good x costs \$2 per unit, good y costs \$1 per unit, and she has \$8 to spend). What does she buy?

The standard technique is to substitute the budget constraint into the utility function and maximise. From the budget we have $y = 8 - 2x$, so we want to maximise $u(x) = x^{½}(8 - 2x)^{½}$

Taking the derivative, setting it to zero, and solving (the algebra is tedious but elementary) gives $x^* = 2$, hence $y^* = 4$. The maximised utility is $u^* = (2 \times 4)^{½} = 2\sqrt{2} \approx 2.83$. Figure 1 shows the picture: three indifference curves — sets of bundles between which the consumer is indifferent — and the budget line. The optimum is where the highest reachable indifference curve just kisses the budget line.

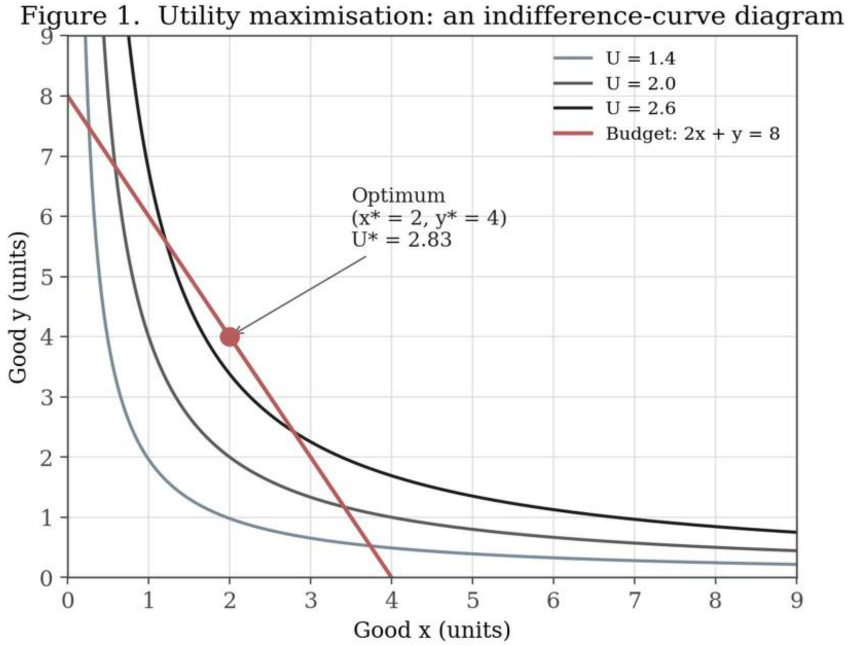

Figure 1. An indifference-curve diagram for a Cobb–Douglas consumer with budget 2x + y = 8.

This is the canonical microeconomic story. The agent's decision problem reduces to optimisation: find the feasible point that maximises utility. But notice what is missing: nowhere in this story does any other agent appear. The prices are fixed exogenously; the consumer does not have to wonder what the seller is planning to do. Game theory begins precisely when we relax this assumption.

Expected utility and risk

Many real choices involve uncertainty. Should you take an umbrella? The answer depends on whether it rains. Should a firm invest in R&D? The answer depends on whether the research succeeds. To handle these, we extend the utility framework to **lotteries**: probability distributions over outcomes. The von Neumann–Morgenstern (vNM) expected utility theorem says that if an agent's preferences over lotteries satisfy a small number of additional axioms (most controversially, the so-called independence axiom), then there exists a utility function u defined on outcomes such that the agent prefers lottery L_1 to lottery L_2 whenever the **expected utility** of L_1 exceeds the expected utility of L_2.

Formally: if lottery L assigns probability pi to outcome xi, then
$EU(L) = \Sigma p_i u(x_i)$

Risk attitudes show up in the shape of u. If u is concave (its slope decreases as outcomes get better), the agent is **risk-averse**: she prefers a sure $50 to a 50-50 gamble between $0 and $100, even though both have expected monetary value $50. If u is linear, she is risk-neutral; if u is convex, she is risk-seeking. Expected utility theory is the workhorse of decision theory under uncertainty, and it is also — crucially — the framework game theorists use to make sense of mixed strategies, randomisation, and uncertainty about other players' types.

From decisions to games

Armed with utility functions, we can now define what a **game** is. A game consists of three things:

- A set of players, usually indexed 1, 2, ..., n.
- For each player, a set of available strategies (which may be pure actions, or randomisations over actions).
- For each combination of strategies — one chosen by each player — a payoff (utility) for each player.

This is the apparatus we will use to build everything that follows. The next chapter introduces our first real game.

Static games and the Prisoner's Dilemma

A **static game** — also called a one-shot game, or a game in *normal form* — is one in which each player chooses a strategy simultaneously, without observing the other players' choices, and the payoffs are realised in a single shot. The simplest such games have only two players with two strategies each, and can be displayed as a 2×2 payoff matrix. The most famous of these is the Prisoner's Dilemma.

The Prisoner's Dilemma

Two suspects, arrested on suspicion of a serious crime, are held in separate cells. The police don't have enough evidence to convict them of the serious crime, but they do have enough to convict each on a minor charge. So the police offer each suspect the same deal: confess (and implicate the other), and you get a reduced sentence — but the more you and your partner cooperate with each other (by staying silent), the worse off you both individually become if you do not also confess.

The standard version of the payoffs runs as follows. If both stay silent ("Cooperate" with each other), they each get 1 year in prison on the minor charge. If both confess ("Defect"), they each get 7 years. If one defects while the other stays silent, the defector walks free (0 years) and the cooperator gets the maximum sentence of 10 years. Following the convention in game theory, we write payoffs as utilities — here, negative years in prison, since fewer years is better. The result is the matrix in Figure 2.

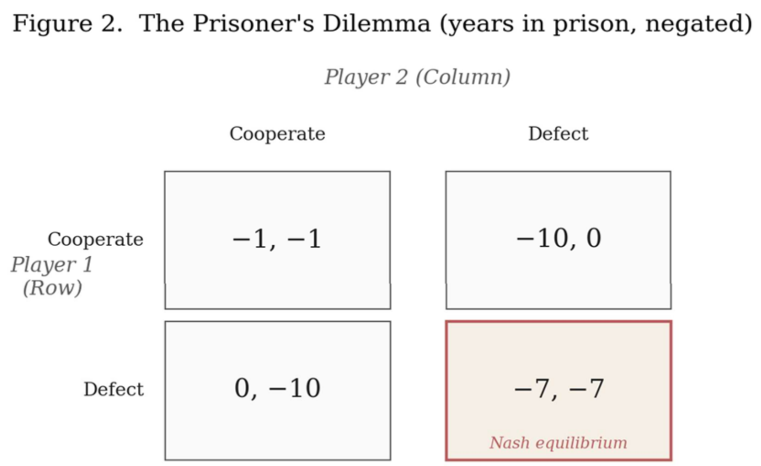

Figure 2. Payoff matrix for the Prisoner's Dilemma. Each cell shows (Row's payoff, Column's payoff).

Dominant strategies

To analyse this game, we ask a simple question from each player's perspective: regardless of what the other player does, what is my best move?

Take Player 1 (Row). If Player 2 cooperates, Row's choices are −1 (cooperate) or 0 (defect). Defecting is better. If Player 2 defects, Row's choices are −10 (cooperate) or −7 (defect). Defecting is still better. So no matter what Player 2 does, Row prefers to defect. Defection is a **dominant strategy** for Row. By the symmetry of the payoffs, the same is true for Column. Both players defect.

Dominant strategy. A strategy is (strictly) dominant for a player if it yields a strictly higher payoff than every other strategy, for every possible choice by the other players. When a dominant strategy exists, a rational player will always play it.

The outcome — both defect, each gets 7 years — is the famous **paradox** of the Prisoner's Dilemma. Both players, acting in their narrowly defined self-interest, end up in an outcome that is strictly worse for both of them than the cooperative outcome (where each would have served only one year). The pursuit of individual rationality has led to collective irrationality. This is not a mathematical curiosity but a structural feature of an enormous range of real-world problems, from arms races to overfishing to climate change to public goods provision in general.

Nash equilibrium

Dominant strategies are wonderfully clean when they exist, but most games don't have them. We need a more general solution concept. That concept is John Nash's: a **Nash equilibrium** is a combination of strategies — one for each player — such that no player can improve her payoff by unilaterally changing her own strategy, given what the others are doing.

Nash equilibrium. A strategy profile $(s_1^*, s_2^*, ..., s_n^*)$ is a Nash equilibrium if, for every player i and every alternative strategy s_i, $u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$ — where s_{-i}^* denotes the equilibrium strategies of all players other than i. In words: at a Nash equilibrium, each player is best-responding to the others.

In the Prisoner's Dilemma, (Defect, Defect) is the unique Nash equilibrium. If Row plays Defect and Column considers switching to Cooperate, her payoff drops from −7 to −10. So Cooperate is not a profitable deviation. Symmetrically for Row. Neither player can do better by deviating, so the profile is stable.

Notice that (Cooperate, Cooperate), although it is **Pareto-superior** — both players are strictly better off there — is *not* a Nash equilibrium, because each player has an incentive to defect unilaterally. This gap between what's individually optimal and what's collectively optimal is one of the most important conceptual insights in twentieth-century social science.

Existence: Nash's theorem

Does every game have a Nash equilibrium? In pure strategies — where each player simply picks one action — the answer is no. The children's game Rock–Paper–Scissors has no pure-strategy equilibrium: whatever a player does, her opponent can do strictly better by choosing the action that beats it. But if we allow players to **randomise** over their actions, the answer becomes yes.

Nash's existence theorem (1950). Every finite normal-form game has at least one Nash equilibrium, possibly in mixed strategies.

The proof uses a deep result in mathematics called Brouwer's fixed-point theorem, which is well beyond the scope of this paper. But the conclusion is striking: there is always at least one stable strategic configuration. The remaining question is what it looks like, and how to find it.

III. Mixed Strategies And Randomisation

Some games — like Rock–Paper–Scissors, or a penalty kick in football — have no pure-strategy Nash equilibrium. The reason is structural: in these games, each player's best action depends on what the other will do, but if both players' best actions are deterministic, they form an inconsistent loop. The way out is randomisation.

Matching Pennies

Consider a stripped-down version of Rock–Paper–Scissors, called **Matching Pennies**. Each of two players simultaneously chooses Heads or Tails. If the choices match, Row wins $1 from Column; if they differ, Column wins $1 from Row. The payoff matrix is:

		Column	
		Heads	*Tails*
Row	*Heads*	+1, −1	−1, +1
	Tails	−1, +1	+1, −1

There is no pure-strategy equilibrium. To see this, suppose Row plays Heads with certainty. Then Column's best response is Tails (winning $1 instead of losing $1). But if Column plays Tails, Row's best response is Tails too, not Heads. Whatever pure strategy we start with, somebody wants to deviate.

Solving for the mixed-strategy equilibrium

The standard trick is to make each player indifferent between her two pure actions, by choosing the right mixing probability for the *other* player. Let Column play Heads with probability q (and Tails with probability 1−q). Then Row's expected payoffs from each of her pure strategies are:

Worked example: Row's indifference condition (Matching Pennies) If Row plays Heads: $EU = (+1)\cdot q + (-1)\cdot(1-q) = 2q - 1$ If Row plays Tails: $EU = (-1)\cdot q + (+1)\cdot(1-q) = 1 - 2q$ Setting them equal: $2q - 1 = 1 - 2q \rightarrow 4q = 2 \rightarrow q = 1/2$

So if Column plays Heads and Tails with equal probability, Row is indifferent between her two pure strategies — and therefore willing to mix. By the symmetry of the game, the same calculation shows that Row must also mix 50–50 to keep Column indifferent. The unique Nash equilibrium is for each player to randomise uniformly, with expected payoff zero. Figure 3 displays this as the intersection of two best-response correspondences.

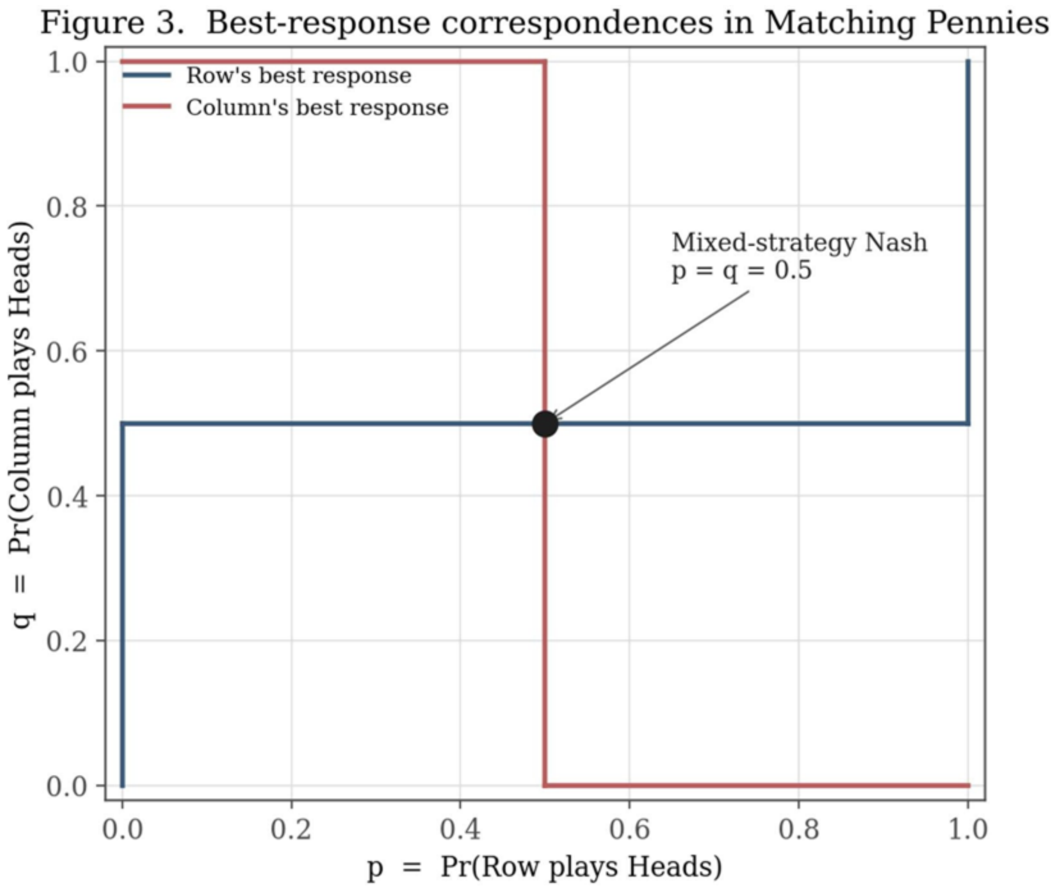

Figure 3. Best-response correspondences for Matching Pennies. The blue and red curves cross at p = q = 0.5, the mixed-strategy Nash equilibrium.

Why randomisation makes sense

Mixed strategies sometimes strike newcomers as artificial. Does a penalty taker really compute a probability distribution before kicking? Probably not consciously. But the prediction the model makes — that, over many penalties, kickers will appear to randomise between left, right, and centre in a way that keeps goalkeepers indifferent — is borne out astonishingly well in empirical data. Ignacio Palacios-Huerta's analyses of professional football penalties show kickers and keepers playing almost exactly the mixed strategy theory predicts, even though no one running onto the pitch is doing the algebra.

A useful interpretation is that the mixed strategy is not in any individual's head but in the population: across many encounters, the empirical frequency of choices traces out the mixing probabilities. The strategies are evolutionary, not deliberate. The mathematics still tells us where the system ends up.

IV. Coordination Games And Multiple Equilibria

Not all interactions are conflicts. Many involve coordination: drivers want to agree on which side of the road to drive on, software users want to use the same file formats, and friends want to meet at the same restaurant. These coordination games typically have multiple Nash equilibria, which raises the question: how, in practice, do real players select one?

Battle of the Sexes

Two friends want to spend the evening together. One prefers a concert; the other prefers a film. But both prefer being together (whatever the activity) to being apart. The payoff matrix is:

		Sam	
		Concert	*Film*
Alex	*Concert*	2, 1	0, 0
	Film	0, 0	1, 2

There are two pure-strategy Nash equilibria: (Concert, Concert) and (Film, Film). At each, neither player can improve by unilaterally switching, since switching means going alone and getting 0. There is also a mixed-strategy equilibrium. To find it, suppose Sam plays Concert with probability q. Alex is indifferent when $2q + 0\cdot(1-q) = 0\cdot q + 1\cdot(1-q)$, i.e. $q = 1/3$. Symmetrically, Alex must play Concert with probability 2/3. The mixed equilibrium yields each player an expected payoff of 2/3, which is worse than either pure equilibrium — a frustrating outcome that highlights the cost of failed coordination.

The Stag Hunt

A more profound coordination problem is the Stag Hunt, due originally to Jean-Jacques Rousseau. Two hunters can each choose to hunt a stag (which requires cooperation: neither can catch one alone) or a hare (which can be caught alone, but is much smaller). The payoffs are:

		Hunter 2	
		Stag	*Hare*
Hunter 1	*Stag*	5, 5	0, 3
	Hare	3, 0	3, 3

There are again two pure-strategy equilibria: (Stag, Stag), which gives each hunter 5, and (Hare, Hare), which gives each hunter 3. The first is **payoff-dominant** — both players strictly prefer it. But the second is **risk-dominant**: it requires no trust in the other player. Hunting the stag is wonderful if your partner does too, but disastrous (payoff 0) if she runs after a hare. Hunting the hare is safe whatever your partner does.

How risky is Stag, exactly? Suppose Hunter 1 believes Hunter 2 will hunt Stag with probability p. Then her expected payoffs are $EU(Stag) = 5p$ and $EU(Hare) = 3$. She is indifferent at $p = 3/5 = 0.6$. So she should hunt the stag only if she is at least 60% confident her partner will too. Figure 6 visualises this: the more pessimistic you are about your partner, the more attractive the safe option becomes.

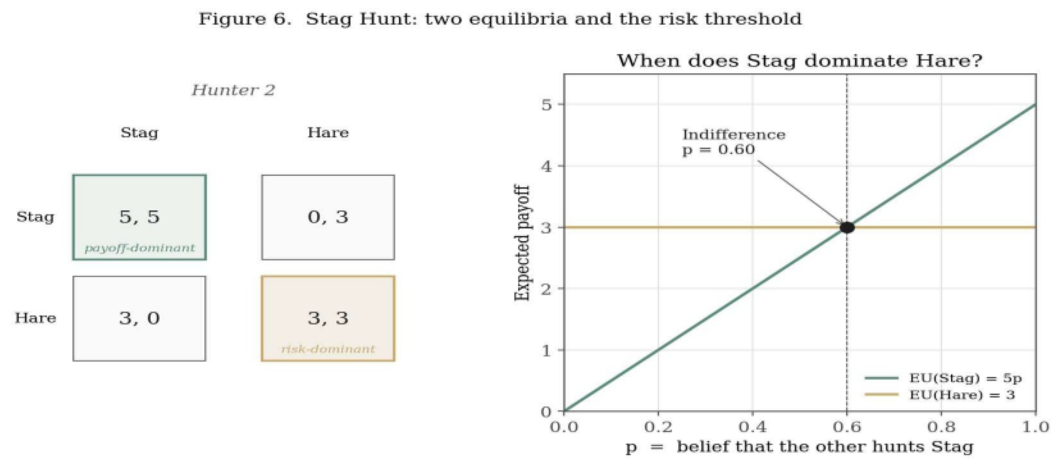

Figure 6. The Stag Hunt has two pure-strategy equilibria. The left panel marks them on the matrix; the right panel plots the expected payoff to each pure strategy as a function of belief about the partner.

The Stag Hunt has been used to model an enormous variety of social phenomena: the founding of social contracts, the formation of cartels, the adoption of new technologies. In each case the structural insight is the same — the cooperative outcome is better for everyone, but only if everyone is sufficiently confident that everyone else will cooperate too. Institutions, communication, repeated interaction, and reputation are all mechanisms human societies have evolved to push the equilibrium from the risk-dominant (Hare) to the payoff-dominant (Stag).

V. Continuous Strategies: The Cournot Duopoly

So far our games have had finite strategy sets (Cooperate / Defect, Stag / Hare). But in many economic settings, the strategy is a continuous quantity — how much output a firm produces, what price to charge, how much to invest. The same Nash apparatus applies; we simply use calculus instead of comparing cells in a matrix. The canonical example is Augustin Cournot's 1838 model of duopoly, which predates the formal apparatus of game theory by more than a century and is now recognised as one of its earliest applications.

Setting up the model

Two firms produce identical units of a good. Let q_1 and q_2 be the quantities produced by firms 1 and 2 respectively. The market price is determined by the inverse demand function

$P = a - b(q_1 + q_2)$

where $a > 0$ and $b > 0$ are parameters. Each firm has a constant marginal cost c, and (for simplicity) no fixed costs. Firm i's profit is therefore

$\pi_i(q_i, q_j) = (P - c) \cdot q_i = [a - c - b(q_i + q_j)] \cdot q_i$

Each firm chooses its quantity simultaneously, taking the other's quantity as given. We want to find the Nash equilibrium pair (q_1^*, q_2^*).

Each firm's best response

Take firm 1's perspective. Treating q_2 as fixed, it maximises π_1 by setting $\partial\pi_1/\partial q_1 = 0$:

Worked example: Firm 1's first-order condition

$\pi_1 = (a - c - b\, q_1 - b\, q_2) \cdot q_1$

$= (a - c)\, q_1 - b\, q_1^2 - b\, q_2\, q_1$

$\partial\pi_1/\partial q_1 = (a - c) - 2\, b\, q_1 - b\, q_2 = 0$

$\Rightarrow q_1 = (a - c - b\, q_2) / (2b)$ ← Firm 1's reaction function

By symmetry, firm 2's reaction function is $q_2 = (a - c - b\, q_1) / (2b)$. These two equations describe each firm's best response to the other's output. The Nash equilibrium is where they intersect: both firms are simultaneously best-responding.

Solving for the equilibrium

Setting $q_1 = q_2 = q^*$ (by symmetry the equilibrium must be symmetric) and substituting:

Worked example: Cournot-Nash equilibrium

$q^* = (a - c - b\, q^*) / (2b)$

$2b\, q^* = a - c - b\, q^*$ $3b\, q^* = a - c$

$q^* = (a - c) / (3b)$ ← each firm's equilibrium output Total $Q^* = 2\, q^* = 2(a - c) / (3b)$

$P^* = a - b\, Q^* = (a + 2c) / 3$

$\pi_i^* = (P^* - c)\, q^* = (a - c)^2 / (9b)$

Let us put numbers on this. Suppose a = 100, b = 1, c = 20 (units of the good, dollars). Then each firm produces $q^* = 80/3 \approx 26.67$ units. Total industry output is about 53.33 units. The price is $(100 + 40)/3 \approx \$46.67$. Each firm earns profit $(46.67 - 20) \times 26.67 \approx \$711$.

Figure 4. Cournot duopoly: reaction curves intersect at Nash equilibrium

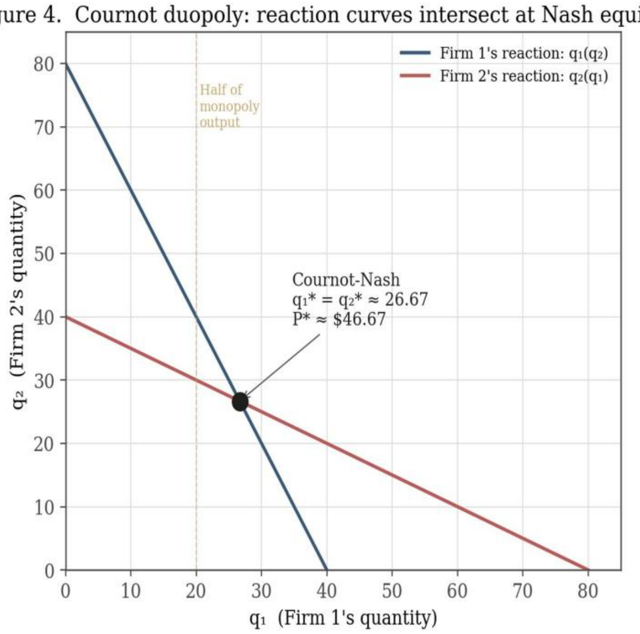

Figure 4. Reaction curves for the Cournot duopoly with a = 100, b = 1, c = 20. The Nash equilibrium lies at their intersection, $q_1 = q_2 \approx 26.67$.

Comparison with monopoly and perfect competition

It is illuminating to compare the Cournot duopoly equilibrium with the outcomes a single monopolist or a perfectly competitive market would produce in the same setting:

Market structure	Total Q	Price P	Industry profit
Perfect competition	80	$20	$0
Cournot duopoly	53.3	$46.67	$1,422
Monopoly	40	$60	$1,600

The Cournot duopoly sits between the two extremes. Output is lower than in perfect competition (so the price is higher and consumers worse off), but higher than under monopoly (so industry profit is lower). This is a remarkable result: even with only two firms competing, strategic interaction pushes them part of the way toward the competitive outcome — but not all the way, because each firm restrains its output to keep the price up. The deeper insight is that the competitive equilibrium is the limiting case as the number of Cournot firms goes to infinity.

VI. Dynamic Games And Backward Induction

Static games assume simultaneous, one-shot moves. But many strategic situations are **dynamic**: players move in sequence, observing what has happened so far before they choose. A firm decides whether to enter a market; the incumbent then decides whether to fight or accommodate. A union decides whether to strike; management then decides whether to negotiate or lock out. The natural representation of such games is the **extensive form**: a tree of decision nodes, with branches for each available action.

Extensive form: an entry-deterrence example

Figure 5 shows a classic entry-deterrence game. A potential Entrant decides whether to Enter the market or Stay out. If she stays out, the Incumbent enjoys a monopoly: payoffs (0, 10). If she enters, the Incumbent must choose between Fighting (a costly price war) or Accommodating (sharing the market). Fight leaves both firms badly off: payoffs (−3, 1). Accommodate splits the market: payoffs (4, 4).

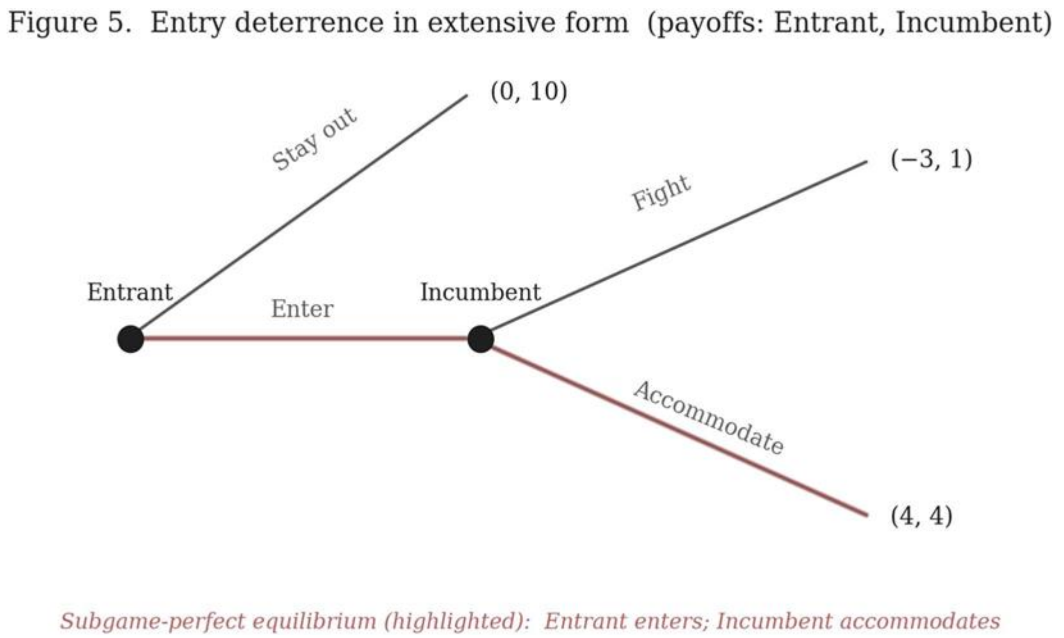

Figure 5. Entry deterrence as an extensive-form game. The pink path shows the subgame-perfect equilibrium: the Entrant enters and the Incumbent accommodates.

Backward induction

To find the equilibrium, we reason backwards from the end of the tree. Suppose the Entrant has entered. The Incumbent must choose: Fight gives a payoff of 1, Accommodate gives 4. The Incumbent will Accommodate. Now move one step back. The Entrant, anticipating that the Incumbent will Accommodate if she enters, compares Enter (payoff 4) with Stay out (payoff 0). She enters. The equilibrium is (Enter, Accommodate), with payoffs (4, 4).

Backward induction. In a finite extensive-form game of perfect information, the unique rational solution can be found by starting at the terminal nodes, identifying each player's optimal action at the last decision, then folding the tree back one step at a time. Each player at each decision node assumes that all future decisions will be played optimally.

Subgame perfection and incredible threats

Now suppose the Incumbent announces, before the game starts: "If you enter, I will Fight you to the death." If the Entrant believed this, she would stay out, and the Incumbent would enjoy the monopoly payoff of 10. So the threat would be effective — but is it credible?

The answer is no. Once the Entrant has actually entered, the Incumbent's best move is to Accommodate (payoff 4) rather than Fight (payoff 1). The threat to fight is *ex post* irrational. Game theorists call such threats **incredible**: they would not survive the moment of truth. The Nash equilibrium concept on its own can't rule them out — there is in fact a Nash equilibrium in which the Entrant Stays out and the Incumbent (counterfactually) plans to Fight. But this equilibrium relies on an irrational off-path threat, and the refinement called **subgame-perfect equilibrium** — which requires that strategies form a Nash equilibrium in every subgame — eliminates it.

This insight has enormous practical bite. In international trade, governments often threaten retaliation to deter unfavourable behaviour. Whether such threats are credible depends on whether the government would actually carry them out once provoked. Building credibility — through reputation, repeated play, or institutional commitments ("tying one's hands") — is one of the central problems in political economy.

VII. Repeated Games And The Emergence Of Cooperation

Return to the Prisoner's Dilemma. In a one-shot game, both players defect. But what if they play the same game repeatedly? What if there is a tomorrow in which my defection today can be punished? Repeated games — also called supergames — radically change what behaviour is rational.

Finite versus infinite repetition

If the Prisoner's Dilemma is played a known, finite number of times, backward induction unravels cooperation. In the last round, both players defect (as in the one-shot game). Knowing that, both players defect in the second-to-last round (there is no future to lose by defecting). The argument cascades all the way back to round 1. The unique subgame-perfect equilibrium is to defect every round.

This is a surprisingly grim result. The intuition is that without a future, there is no shadow of punishment. But the picture changes dramatically if the game has no fixed endpoint — if there is always some positive probability that the players will meet again. In an infinitely repeated game, cooperation can be sustained.

The folk theorem

Suppose players discount future payoffs by a factor $\delta \in (0, 1)$ per period, so that one unit of utility next period is worth δ today. Consider the following strategy for the Prisoner's Dilemma: "Cooperate, but if my opponent ever defects, defect forever after." This is the **Grim Trigger** strategy.

Is it a Nash equilibrium for both players to play Grim Trigger? On the equilibrium path, both cooperate forever, yielding a present discounted value of payoffs:

$V_C = 3 + 3\delta + 3\delta^2 + ... = 3 / (1 - \delta)$

Now suppose one player deviates and defects in some period. She earns 5 in that period (the temptation payoff against a cooperator), but her opponent will defect forever after. So her payoff from defection is

$V_D = 5 + \delta \cdot 1 + \delta^2 \cdot 1 + ... = 5 + \delta / (1 - \delta)$

Cooperation is sustainable when $V_C \geq V_D$, i.e.

Worked example: When does Grim Trigger sustain cooperation?

$3 / (1 - \delta) \geq 5 + \delta / (1 - \delta)$

$3 \geq 5(1 - \delta) + \delta$

$3 \geq 5 - 4\delta$

$\delta \geq 1/2$

So if the discount factor is at least 1/2 — meaning players value the future enough — cooperation is rational. This is a special case of the celebrated **Folk Theorem**, which says that in infinitely repeated games with sufficiently patient players, any outcome in which each player gets at least her minimax payoff can be sustained as a subgame-perfect equilibrium. There are many such equilibria; the theory does not pin down a unique outcome.

Axelrod's tournament: tit-for-tat

In 1980 the political scientist Robert Axelrod ran a celebrated computer tournament. He invited researchers to submit strategies for playing 200-round Iterated Prisoner's Dilemma matches; each strategy played a round-robin against all others. The winning strategy was submitted by the mathematician Anatol Rapoport and consisted of just four lines of code. It is called **Tit-for-Tat**:

- On the first move, cooperate.
- On every subsequent move, do whatever the opponent did on the previous move.

Tit-for-Tat is nice (it never defects first), retaliatory (it punishes defection immediately), forgiving (a single cooperative move from the opponent gets it back to cooperating), and clear (its rule is transparent and predictable). Figure 7 shows a small replication of this kind of tournament. Tit-for-Tat does not *beat* anyone in head-to-head play — at best it ties — but it accumulates the highest total score across all matchups, because it never gets exploited badly and reliably reaches the (Cooperate, Cooperate) outcome with other cooperative strategies.

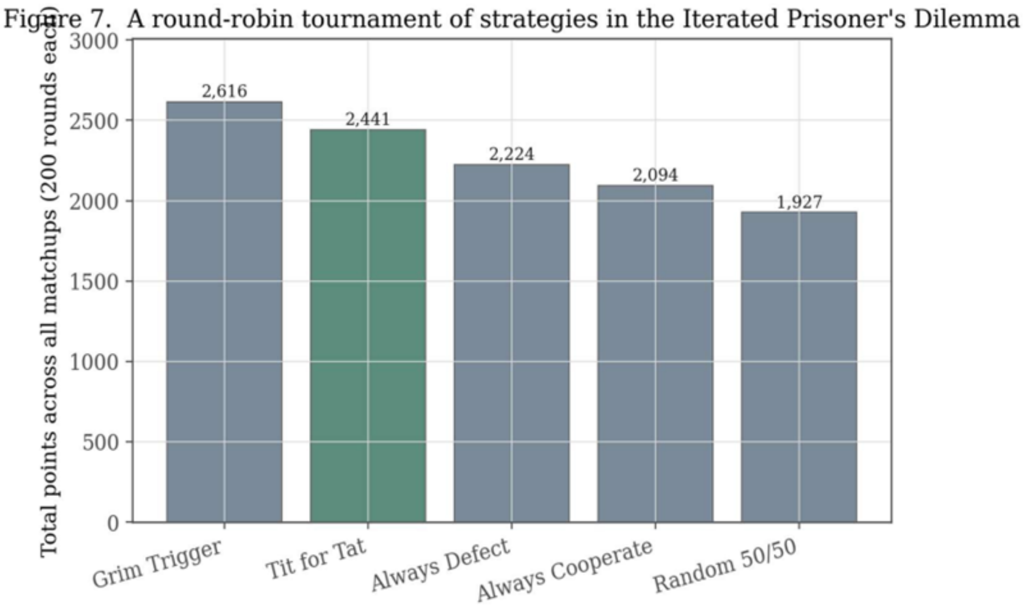

Figure 7. Total tournament scores for five strategies in a 200-round Iterated Prisoner's Dilemma round-robin. Tit-for-Tat (green, highlighted) finishes near the top — beaten only narrowly here by Grim Trigger, because the unconditional retaliation of Grim avoids the costly alternating cycle that Tit-for-Tat falls into when matched against Always Defect.

In the original 1980 tournament against fifteen rival strategies, Tit-for-Tat won outright; in a wide range of subsequent replications it remains either the winner or a top-three finisher. The deeper lesson, though, is independent of which conditional cooperator happens to come first in any one simulation. Cooperation is not a mystery or an irrational deviation from self-interest. In any repeated interaction with a long-enough shadow of the future, conditional cooperation — punishing defection, rewarding cooperation — is a rational strategy. This insight has illuminated everything from the evolution of altruism in biology to the stability of international agreements to the operation of trust networks in informal trade.

VIII. Asymmetric Information And The Market For Lemons

Almost everything we've considered so far has assumed **complete information**: each player knows the payoff structure of the game, the strategies available, and the rationality of the other players. In reality, players often know things their opponents don't. A used-car seller knows whether her car is a lemon; the buyer does not. A job applicant knows her ability; the employer does not. A borrower knows whether she's likely to default; the bank does not. These information asymmetries can have spectacular consequences.

Akerlof's lemons

In a 1970 paper that would eventually win him a Nobel Prize, George Akerlof showed that asymmetric information can cause markets to **unravel** — to shrink, or even disappear, despite the existence of mutually beneficial trades. His example was the used-car market. Suppose every used car has a true quality v, distributed uniformly on the interval [0, V] (so V is the best possible car, and the average car has quality V/2). The seller knows v; the buyer does not. Buyers, who value cars more highly than sellers do (because they want to use them more), are willing to pay k · v for a car of quality v, where k > 1 captures the gains from trade.

In a market with full information, every car would trade: for each car worth v to the seller, the buyer is willing to pay kv > v. But under asymmetric information, the buyer doesn't know v. She knows only the price P being asked. Suppose all cars with v ≤ P are offered for sale (since for those cars, P meets or beats the seller's reservation value). The average quality of cars offered is then P/2. The buyer is willing to pay at most k · (P/2) for an average car of unknown quality. For the market to function, we need

$$k \cdot P/2 \geq P \quad \Leftrightarrow \quad k \geq 2$$

If the gains from trade aren't large enough — if k < 2 — the buyer's willingness to pay is always less than the price asked. The only price at which the market clears is P = 0: the market collapses entirely, even though for the highest-quality cars, both parties would have been delighted to trade at a price slightly above v. Figure 10 illustrates this with k = 1.5.

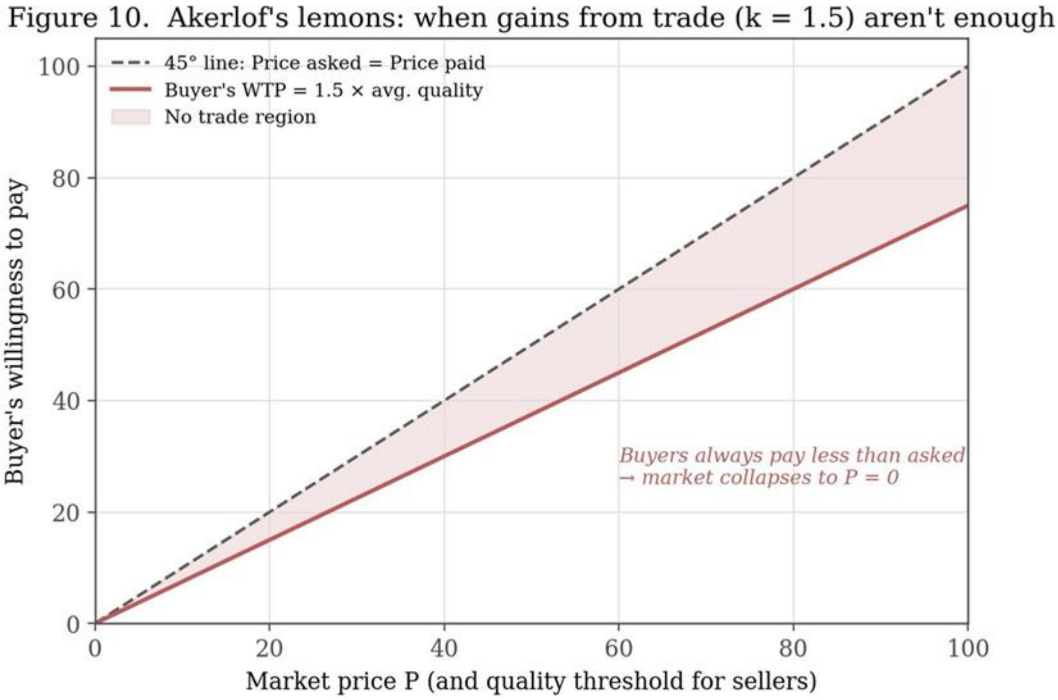

Figure 10. Akerlof's lemons. With k = 1.5, the buyer's willingness to pay (red) is always below the price (dashed line), and the market unravels to P = 0.

Signalling and screening

Akerlof's result is so devastating that one might expect used-car markets, insurance markets, and credit markets to barely exist. They do exist, of course — and the reason is that participants develop **signalling** and **screening** mechanisms to combat the information problem.

Michael Spence's 1973 job-market signalling model — also Nobel-recognised — explains why workers pursue costly education even if education itself adds nothing to their productivity. If education is harder for low-productivity workers than high-productivity ones, then a sufficiently demanding degree can serve as a signal of ability, and a **separating equilibrium** emerges: high-types invest in the signal, low-types do not, and the labour market sorts workers efficiently despite asymmetric information.

Joseph Stiglitz's screening models flip the logic: it is the uninformed party (insurer, employer, lender) who designs a menu of contracts such that different types of agents self-select. A high-deductible insurance plan with low premiums attracts safe drivers; a low-deductible plan with high premiums attracts riskier ones. Together, signalling and screening form the bedrock of modern contract theory.

IX. Auctions And Mechanism Design

Auctions are a natural application of game theory: each bidder must decide how much to bid, knowing that her optimal bid depends on what the others are likely to bid, which depends in turn on how much the others value the item. Auction theory studies how different auction rules — first-price, second-price, English, Dutch — affect equilibrium behaviour and the seller's expected revenue.

First-price and second-price sealed-bid auctions

In a **first-price sealed-bid** auction, each of n bidders submits a sealed bid. The highest bidder wins and pays her bid. In a **second-price sealed-bid** auction (also called a Vickrey auction), the highest bidder wins but pays the second-highest bid. The strategic implications are very different.

In a Vickrey auction it is a **weakly dominant strategy** to bid your true value. Why? Suppose your true value is v_i. If you bid more than v_i, you only change the outcome by winning auctions you'd otherwise lose — and in those cases, the second-highest bid (which determines the price you'd pay) might be more than v_i, leaving you worse off. If you bid less than v_i, you only change the outcome by losing auctions you'd otherwise win — and you'd be giving up surplus. Truthful bidding is optimal. This is one of the cleanest results in mechanism design.

Bid shading and the FPSB equilibrium

In a first-price auction, by contrast, you must **shade** your bid below your value to leave yourself some surplus when you win. How much to shade depends on what others are likely to bid. Assume bidders' valuations are independent draws from the uniform distribution on [0, 1]. Then it can be shown that the symmetric Bayes-Nash equilibrium bid for a bidder with value v_i is

$b_i^*(v_i) = (n - 1)/n \cdot v_i$

With two bidders, each shades to half her value. With ten bidders, each shades to 90%. As competition intensifies, bidders shade less because the chance of being outbid grows.

The revenue equivalence theorem

One of the most beautiful results in auction theory, due to William Vickrey and generalised by Roger Myerson, is the **revenue equivalence theorem**: under fairly weak conditions (independent private values, risk-neutral bidders, symmetry), all auction formats that allocate the object to the highest-valuing bidder and give zero surplus to a bidder with the lowest possible value yield the same expected revenue.

Concretely, with n bidders and uniform[0, 1] values, the expected revenue under either auction format is $E[\text{revenue}] = (n - 1) / (n + 1)$

This is exactly the expected value of the second-highest order statistic from n draws of a uniform distribution — which is the price the Vickrey auction extracts, and also (by the revenue equivalence) the expected payment under first-price auctions. Figure 9 plots both quantities as a function of the number of bidders.

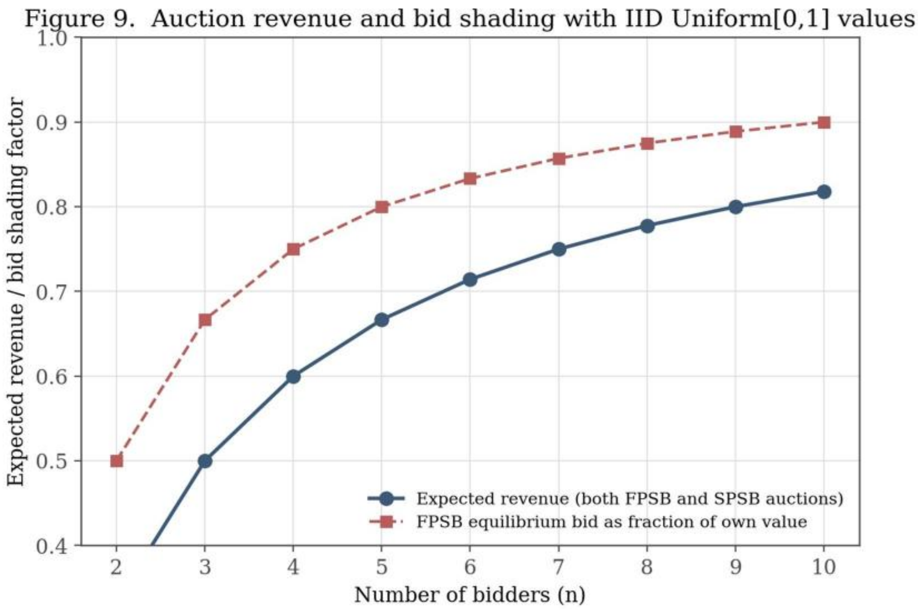

Figure 9. Expected revenue and FPSB bid shading factor in a Uniform[0, 1] private-values auction with n bidders.

The revenue equivalence theorem is a foundational result for **mechanism design** — the engineering branch of game theory which asks: given that we want a certain outcome (say, allocating spectrum licences to the firms that value them most), what rules should we set? Mechanism design has led, among other things, to the multi-billion-dollar spectrum auctions run by governments around the world, and is the field for which Alvin Roth and Lloyd Shapley shared the 2012 Nobel Prize.

X. When The Theory Meets The World

The mathematical apparatus we have built is elegant. Whether it is *descriptive* of how real people actually behave is another question. Beginning in the 1980s, experimental economists and behavioural game theorists have tested the predictions of classical game theory in the laboratory and in the field, often with surprising results.

The ultimatum game

The simplest and most striking example is the **ultimatum game**, introduced by Werner Güth, Rolf Schmittberger and Bernd Schwarze in 1982. Two players are randomly paired. One — the proposer — is given a sum of money (say $10) and must offer some share of it to the other — the responder. The responder either accepts (each player gets the proposed split) or rejects (both get zero). The game is played once, anonymously.

The textbook prediction is stark. The responder, faced with any positive offer, should accept (some money is better than none). Anticipating this, the proposer should offer the smallest possible positive amount, say 1 cent. The unique subgame-perfect equilibrium is a near-zero offer that is accepted.

Real human beings do not play this equilibrium. The modal offer in industrialised-country samples is 40–50% of the pie. Offers below 20% are rejected over half the time. People appear to care not just about their own monetary payoff but about fairness, equality, and what philosophers call "the social emotions" — anger at being treated unfairly, guilt at treating someone else unfairly.

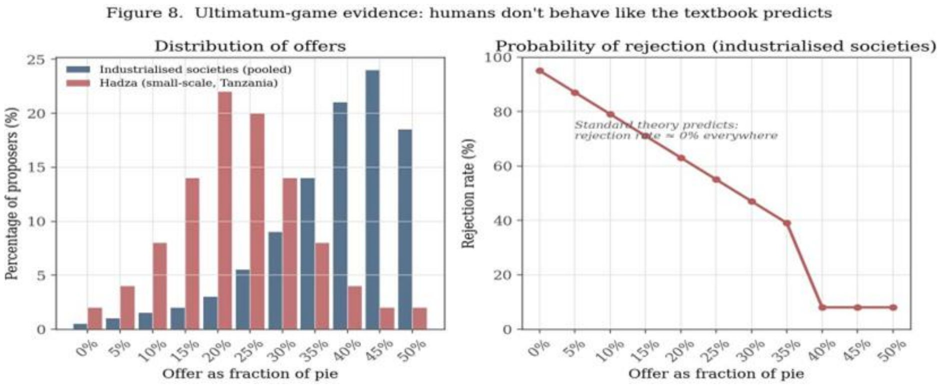

Figure 8. Stylised distributions of offers and rejection rates in ultimatum-game experiments. Theory predicts near-zero offers always accepted; experiments find generous offers and substantial rejection of low ones.

Cross-cultural evidence

Are these behaviours universal, or culturally specific? A landmark study by Joseph Henrich and colleagues, published in 2001, ran ultimatum-game and related experiments in fifteen small-scale societies, ranging from foraging bands in Tanzania (the Hadza) and the Peruvian Amazon (the Machiguenga) to whale-hunting villages in Indonesia (Lamalera) to nomadic herders in Mongolia. The variation was enormous.

- The Machiguenga of Peru, a relatively isolated horticultural society with little market exchange, made modal offers of 15% and accepted almost everything.
- The Lamalera, who hunt whales cooperatively and share the catch by elaborate rules, made offers above 50% — they actively wanted to give more than half away.
- The Hadza, a foraging society in Tanzania, made low offers (around 20%) but also rejected low offers frequently, displaying conflict between giving little and punishing those who gave little.

The variance across societies was several times larger than the variance within any one society, and was strongly correlated with two ecological variables: the importance of market integration in the local economy, and the payoff to cooperation in everyday productive activities. Where everyday life rewarded cooperation, ultimatum-game players cooperated. Where it did not, they did not. The textbook "selfish rational agent" was conspicuously rare anywhere.

What does this mean for game theory?

There are three ways economists have responded to this kind of evidence.

The first response is **behavioural**: modify the utility function to include considerations beyond own monetary payoff. Models like the Fehr–Schmidt inequity-aversion model assume players dislike both being worse off than others (envy) and being better off than others (guilt), in calibrated proportions. With such utility functions, the observed ultimatum-game behaviour becomes a Nash equilibrium of a slightly different game. This is now standard practice in the field.

The second response is **evolutionary**: treat strategies not as the products of conscious optimisation but as cultural norms that have been selected for working well in the small-scale, repeated interactions that characterised most of human history. In a small society where everyone knows everyone else, accepting humiliating offers and never punishing cheats is a recipe for being exploited indefinitely. The ultimatum-game behaviour we observe in the lab may be a residue of strategies that were once perfectly rational.

The third response is **methodological**: classical game theory is a normative theory about what fully rational agents would do, not a descriptive theory about what actual humans do. Both have a place. The mathematics tells us the structure of strategic problems and the long-run gravitational pull of incentives; experimental and behavioural work tells us what people actually do in the short run and what frictions matter. Modern economics increasingly tries to do both at once.

XI. Conclusion

This paper has tried to introduce game theory not as a collection of clever puzzles but as the working microfoundation of modern economics. Its arc has been from the foundations — preference, utility, the rational chooser — through the central solution concept of Nash equilibrium, into the rich worlds of dynamic strategy, repeated interaction, asymmetric information, and market design, and finally up against the messy, illuminating record of how real human beings actually play.

What game theory has taught economics

Three big lessons have emerged from this trajectory.

The first is that **individual rationality and collective rationality can pull in opposite directions**. The Prisoner's Dilemma is the canonical illustration, but the structure recurs everywhere from arms races to climate negotiations to free-riding on public goods. Economic outcomes are not simply the sum of individually optimal decisions; they depend on the architecture of incentives, and that architecture is something institutions can — and do — engineer.

The second is that **information matters as much as preferences**. The lemons model and its descendants showed that markets can fail not because anyone is irrational, but because of who knows what and when. Modern economics is to a large extent the economics of information: who has it, who lacks it, what mechanisms emerge to bridge the gap.

The third is that **the future shapes the present**. In repeated and dynamic games, the shadow of future interaction can sustain cooperation that would be impossible in a one-shot encounter. Reputation, retaliation, conditional kindness — these are not mysterious moral residues but rational responses to repeated play, and they help explain how cooperation can flourish in a world of self-interested agents.

Limitations and open questions

The theory has limitations. Equilibrium concepts often fail to single out a unique prediction — coordination games have multiple equilibria, the Folk Theorem produces a continuum of them. Computing equilibria in large games is mathematically intractable. And, as the experimental evidence makes clear, real humans are not the utility-maximising agents the basic theory assumes. They are emotional, cognitively limited, social, and embedded in cultural norms.

These limitations have given rise to thriving subfields: equilibrium refinement, computational game theory, behavioural and experimental economics, evolutionary game theory. Each is a partial response to a real shortcoming of the classical apparatus. None has displaced it. Game theory, like the rest of economics, advances by accumulation: new layers added to old foundations, with the foundations themselves periodically inspected and reinforced.

Why this matters

For a high-school student standing at the edge of this field, the most important thing to take away is perhaps not any specific theorem but a way of seeing. Strategic interaction is everywhere: in markets and politics, in friendships and rivalries, in algorithms and ecosystems. Game theory provides a precise language for thinking about it. The mathematics can be intimidating, but the underlying ideas

— that each agent's best move depends on what the others are doing, that some outcomes are stable and others are not, that the structure of an interaction shapes who wins and who loses — are not the property of professional economists. They are tools any thoughtful person can use to make sense of the world.

That, in the end, is what makes quantitative economics worth studying. Numbers and equations are not the point. They are the apparatus we use to see clearly. The point is the clarity

References

[1]. Akerlof, G. A. (1970). The Market For "Lemons": Quality Uncertainty And The Market Mechanism. Quarterly Journal Of Economics, 84(3), 488–500.
[2]. Axelrod, R. (1984). The Evolution Of Cooperation. New York: Basic Books.
[3]. Cournot, A. A. (1838). Recherches Sur Les Principes Mathématiques De La Théorie Des Richesses. Paris: Hachette.
[4]. Fehr, E., & Schmidt, K. M. (1999). A Theory Of Fairness, Competition, And Cooperation. Quarterly Journal Of Economics, 114(3), 817–868.
[5]. Güth, W., Schmittberger, R., & Schwarze, B. (1982). An Experimental Analysis Of Ultimatum Bargaining. Journal Of Economic Behavior & Organization, 3(4), 367–388.
[6]. Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., & Mcelreath, R. (2001). In Search Of Homo Economicus: Behavioral Experiments In 15 Small-Scale Societies. American Economic Review, 91(2), 73–78.
[7]. Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). Microeconomic Theory. New York: Oxford University Press.
[8]. Myerson, R. B. (1991). Game Theory: Analysis Of Conflict. Cambridge, MA: Harvard University Press.
[9]. Nash, J. F. (1950). Equilibrium Points In N-Person Games. Proceedings Of The National Academy Of Sciences, 36(1), 48–49.
[10]. Oosterbeek, H., Sloof, R., & Van De Kuilen, G. (2004). Cultural Differences In Ultimatum Game Experiments: Evidence From A Meta-Analysis. Experimental Economics, 7(2), 171–188.
[11]. Osborne, M. J., & Rubinstein, A. (1994). A Course In Game Theory. Cambridge, MA: MIT Press.
[12]. Palacios-Huerta, I. (2003). Professionals Play Minimax. Review Of Economic Studies, 70(2), 395–415. Rousseau, J.-J. (1755). Discours Sur L'origine Et Les Fondements De L'inégalité Parmi Les Hommes. Amsterdam: Marc-Michel Rey.
[13]. Spence, A. M. (1973). Job Market Signaling. Quarterly Journal Of Economics, 87(3), 355–374. Stiglitz, J. E. (1975). The Theory Of "Screening," Education, And The Distribution Of Income. American Economic Review, 65(3), 283–300.
[14]. Vickrey, W. (1961). Counterspeculation, Auctions, And Competitive Sealed Tenders. Journal Of Finance, 16(1), 8–37.
[15]. Von Neumann, J., & Morgenstern, O. (1944). Theory Of Games And Economic Behavior. Princeton, NJ: Princeton University Press.