A Beginner's Tutorial

How to Read *Drakos (2026)* on
LLM Bargaining & Hotel Pricing

Step-by-step through the mathematics
(companion to Drakos (2026), AGEL AI, May 2026)

Audience: someone who knows basic high-school algebra, has not studied advanced mathematics, but wants to understand what the paper says without being scared off by the symbols.
What you will learn: what all the cryptic symbols mean ($\mathbb{E}$, α, $U[0,1]$, $\mathbb{P}$, $\sim$, $\int$), what the Myerson–Satterthwaite theorem says in plain language, why $9/64 \approx 0.844$ is the "best possible" number, what LLMs do when they bargain, and how all of this comes down to a hotel in Rhodes.
How to read it: linearly, with a pencil. In every "Example" box, do the calculation yourself before reading the solution.

Contents

1. Introduction: What is the Problem?

Imagine a hotel in Rhodes and a corporate representative who wants to book 20 rooms for a conference. Both know what they want, but *neither tells the truth to the other*:

- The hotel has in mind a minimum price, say **120 euros/room**, below which it loses money.
- The buyer has a maximum he can pay, say **160 euros/room**.

An agreement is possible. Any price between 120 and 160 euros is better than "no deal" for both parties. But the hotel will not say "120 is my minimum" because then the other side would offer 121. The buyer will not say "160 is my maximum" because then the hotel would demand 159. **Both bluff, and sometimes the deal is lost** even though both parties would have benefited.

Key idea: the central question of the paper

In 1983, Myerson and Satterthwaite proved mathematically that this loss is **unavoidable** when both sides play strategically. Recently, however, when we replace the humans with two **LLMs** (ChatGPT, Claude, Gemini, etc.) and let them bargain, we observe results *above the theoretical bound*. How is that possible? This is what the paper tries to explain.

1.1 What we will learn in this tutorial

1. The language: what the symbols mean.
2. The concept of expected value, the most basic tool.
3. The impossibility theorem: why 9/64 is the bound.
4. How the paper "solves" LLM behaviour theoretically with one parameter α.
5. What happened in practice when 10 models were tested.
6. How all of this applies to a real hotel.

2. The Language of Mathematics: a Dictionary

Before we look at formulas, let us decode the symbols. They are not magic — they are abbreviations.

2.1 Greek letters

Almost all Greek letters are used as names for numbers:

Symbol	Name	In the paper, represents
α	alpha	Probability that the LLM tells the truth (the "central" parameter)
β	beta	Mixture components (binary, continuous, noisy)
δ	delta	A small distance
ϵ	epsilon	Tolerance / error (a very small number)
λ	lambda	Customer arrival rate (per day)
μ, σ	mu, sigma	Mean, standard deviation

Caution: they are not algorithms, they are variables

A Greek letter in mathematics is exactly like x or y in school — just a name for something. If you see $\alpha = 0.7$, it means "the number we call alpha is equal to 0.7".

2.2 The "big" symbols

Symbol	What it means in plain English
$\mathbb{E}[X]$	"The expected value of X" — the average of X if we repeated the experiment many times.
$\mathbb{P}(A)$	"The probability of event A" — a number between 0 (impossible) and 1 (certain).
$X \sim F$	"X follows the distribution F" — it is drawn from some random generator.
$U[0,1]$	"Uniform distribution on $[0,1]$" — every number between 0 and 1 is equally likely.
$\mathbf{1}\{A\}$	"Indicator" — equals 1 if A holds, otherwise 0.
$\int_a^b f(x)\,dx$	"Integral" — the area under the curve of f between a and b.

3. Expected Value: the Most Important Tool

The whole paper relies on expected values throughout. If you understand this, you understand 70% of the paper.

3.1 With a die

You roll a fair die. What do you expect "on average"?

Example: expected value of a die

Each outcome (1, 2, 3, 4, 5, 6) has probability 1/6. The expected value is:

$$\mathbb{E}[X] \;=\; 1\cdot\tfrac{1}{6} + 2\cdot\tfrac{1}{6} + 3\cdot\tfrac{1}{6} + 4\cdot\tfrac{1}{6} + 5\cdot\tfrac{1}{6} + 6\cdot\tfrac{1}{6} \;=\; \tfrac{21}{6} \;=\; 3.5$$

If you roll the die 1000 times and take the average, it will be close to 3.5. You will never get *exactly* 3.5 in a single roll, but this is the value you "expect".

The general rule: expected value = (probability of each outcome) × (the outcome), summed over everything.

3.2 From a die to "any number between 0 and 1"

In the paper, values are not 1, 2, . . . , 6. They are *any number* between 0 and 1. This is called the **uniform distribution**, $U[0,1]$.

Concept: uniform distribution $U[0,1]$

Imagine throwing a dart at a line of length 1, where every point is equally likely. $U[0,1]$ is exactly that.

- Probability of landing in $[0.2, 0.5]$: $0.5 - 0.2 = 0.3$ (30% of the line).
- Probability of landing exactly at 0.5: *zero* (a point has length 0).
- Mean: 0.5 (the middle of the line).

When a distribution is "continuous" (a line, not a die), the expected value is computed with an **integral** instead of a sum. But conceptually it is the same thing.

3.3 Integrals: don't be afraid

Concept: $\int_a^b f(x)\,dx$ in pictures

Sketch the curve $y = f(x)$ on paper. The integral from a to b is the **area** under the curve and above the x-axis, between the vertical lines $x = a$ and $x = b$.
If $f(x) = 1$ (a horizontal line at 1), then from $x = 0$ to $x = 1$ you have a square of side 1, so area 1. Hence $\int_0^1 1\,dx = 1$. No formula needed — just the picture.

Example: a small integral by hand

What is $\int_0^1 x\,dx$?
Solution by picture: draw the line $y = x$ from 0 to 1. It is a diagonal line forming a triangle with the two axes. Base = 1, height = 1, area = $\frac{1}{2} \cdot 1 \cdot 1 = \frac{1}{2}$.
So $\int_0^1 x\,dx = \frac{1}{2}$.
What it means: if you have a random number $X \sim U[0,1]$, its expected value is 0.5 (the middle of the line).

Example: a slightly harder one

What is $\int_0^1 x^2\,dx$?
This is no longer a triangle, it is a curve. You don't need to solve it by hand: the general rule is:

$$\int_a^b x^n\,dx = \frac{b^{n+1} - a^{n+1}}{n+1}$$

So $\int_0^1 x^2\,dx = \frac{1^3 - 0^3}{3} = \frac{1}{3}$.

We will never ask you to solve an integral by hand. But you should know *what it means*: a sum of many small pieces = area.

4. The Bargaining Problem: Sketching the Idea

Now that we know enough, let's see what the paper says. The scene has two protagonists:

- **Buyer**: has a value v in mind (from the English *value*). It is the maximum they will pay.
- **Seller**: has a cost c in mind (from *cost*). It is the minimum they will accept.

In the paper, both are drawn uniformly from $[0, 1]$. So we normalise the amounts: instead of euros, every value is a number between 0 and 1, and all are equally likely.

4.1 When are there "gains from trade"?

If $v > c$, the trade *makes sense*: the buyer values the item more than it costs the seller. The **gain** (or surplus) is:

$$v - c \quad (\text{if } v > c)$$

If $v < c$, the trade is a bad idea for both. This is exactly what equation (1) of the paper says:

$$S(v, c) = (v - c) \cdot \mathbf{1}\{v \geq c\}$$

The $\mathbf{1}\{v \geq c\}$ is the famous "indicator" — equal to 1 if $v \geq c$, else 0.

4.2 The "first-best": what would happen ideally

If a god knew every v and c in the universe, they would say: "trade every time $v > c$". How much average surplus does that yield?

Example: the first-best area

Imagine a 1×1 square where the horizontal axis is c (seller cost) and the vertical axis is v (buyer value). Each point inside the square is a pair "this time the value v and cost c were drawn".
The trade makes sense only when $v \geq c$, i.e. above the diagonal.
The expected surplus is:

$$\mathbb{E}[\text{surplus}] = \int_0^1 \int_0^v (v - c)\, dc\, dv = \frac{1}{6} \approx 0.1667$$

This is the number **1/6** that appears everywhere in the paper. It is called the **first-best efficiency**: the ideal bound.

Important: this $1/6$ is an average when many (v, c) pairs are drawn from $U[0, 1]$. In a specific trade the surplus might be 0, 0.3, 0.7, etc. But *averaged over all possible draws*, it comes out to 0.1667.

4.3 The problem: the buyer bluffs

In reality, the buyer does not tell v to the seller. They state some *lower* number to get a better price. Likewise, the seller states a *higher* number than c. This is called **strategic behaviour**.

Key idea: the Myerson–Satterthwaite problem

When both bluff strategically, some trades are lost. Specifically, the expected surplus drops from $1/6$ (first-best) to:

$$\frac{9}{64} \approx 0.1406$$

Relative to the ideal:

$$\frac{9/64}{1/6} = \frac{9 \cdot 6}{64} = \frac{54}{64} = \frac{27}{32} \approx 0.844 \quad \text{or } 84.4\%$$

So roughly 15.6% of the surplus is lost when they play strategically. And it can be proved that this is the *best any mechanism can do* with strategic players.

5. Why $9/64$? The Chatterjee–Samuelson Equilibrium

Where does this strange number $9/64$ come from? From a formula proposed by Chatterjee–Samuelson in 1983.

5.1 The optimal bluff

If I am a buyer with value v, and I know the seller's cost c is uniform on $[0, 1]$, the mathematically optimal way to bluff is:

$$b^*(v) = \frac{2}{3}v + \frac{1}{12}$$

In words: *I do not offer my true value v, I offer 2/3 of it plus 1/12.*

Example: a numerical bluff

Suppose $v = 0.9$ (I value the item very highly).
My optimal offer is: $b^*(0.9) = \frac{2}{3}(0.9) + \frac{1}{12} = 0.6 + 0.0833 = 0.683$.
So even though I would pay up to 0.9, I offer only 0.683. I keep a "cushion" of 0.217 of own benefit if the deal closes.
If $v = 0.3$ (I value low): $b^*(0.3) = 0.2 + 0.0833 = 0.283$.
Notice that the lower I value, the smaller the bluff — the bluff is not "always -20%", it is a clever function.

Similarly, the seller with cost c will demand:

$$a^*(c) = \frac{2}{3}c + \frac{1}{4}$$

5.2 When does the trade happen with bluffs?

A trade happens only if the buyer's offer $\geq$ the seller's demand:

$$\frac{2}{3}v + \frac{1}{12} \ \geq \ \frac{2}{3}c + \frac{1}{4}$$

Solving for $v - c$:

$$v - c \ \geq \ \frac{1}{4} - \frac{1}{12} = \frac{3-1}{12} = \frac{1}{6}$$

Key idea: the missing gap

So **when they play strategically, the trade happens only if the "true gain" $v - c$ is at least $1/6$ (≈ 0.167).** If v is 0.5 and c is 0.4 (gain 0.1), *the deal is lost* because $0.1 < 1/6$. This is exactly the "δ gap" that appears in the paper in Lemma 6.

6. The Paper's Central Idea: the Parameter α

So far we have seen the classical theory. What does *Drakos (2026)* do differently?

Key idea: the observation that launches the paper

When two LLMs bargain, **they do not play strategically perfectly**. They often tell the truth or something close to it, even when it costs them. So their actual efficiency may be *higher* than the human "best possible".

The paper introduces a simple parameter to measure this:

Concept: the parameter α

$$\alpha = \mathbb{P}(\text{the LLM states its true value when asked})$$

A number between 0 and 1.

- $\alpha = 0$: the LLM always plays strategically (like a human).
- $\alpha = 1$: the LLM always tells the truth.
- $\alpha = 0.7$: tells the truth 70% of the time, bluffs 30%.

α is a property of the model — we measure it, we do not choose it.

6.1 Three kinds of bluffing: when not telling the truth, what does it do?

When the LLM chooses not to tell the truth (probability $1-\alpha$), what does it do? The paper proposes three theoretical cases:

1. **Binary**: either tells the truth completely, or makes the FULL Chatterjee–Samuelson bluff. Like a switch: ON or OFF.
2. **Continuous shading**: a mix. States $\alpha \cdot v + (1-\alpha) \cdot \text{strategic}$. "70% of my answer is truth, 30% is strategic."
3. **Noisy**: states the truth plus a random error. As if you copy a number with a shaky hand.

6.2 The closed-form formulas (central theoretical result)

After some mathematics, the paper arrives at two formulas that let you compute the expected efficiency as a function of α:

Binary:

$$\boxed{E^B(\alpha) \;=\; \frac{9}{64} + \frac{31\alpha}{864} - \frac{17\alpha^2}{1728}}$$

Continuous:

$$\boxed{E^C(\alpha) \;=\; \frac{9(\alpha+1)^2}{8(\alpha+2)^3}}$$

Example: let's check them at the extremes

At $\alpha = 0$ (never truthful):

$$E^B(0) = \frac{9}{64} + 0 - 0 = \frac{9}{64} \approx 0.1406 \checkmark$$

Correct — we recover Chatterjee–Samuelson.
At $\alpha = 1$ (always truthful):

$$E^C(1) = \frac{9 \cdot 4}{8 \cdot 27} = \frac{36}{216} = \frac{1}{6} \checkmark$$

Correct — we hit the first-best.
At a midpoint, $\alpha = 0.5$:

$$E^C(0.5) = \frac{9 \cdot (1.5)^2}{8 \cdot (2.5)^3} = \frac{9 \cdot 2.25}{8 \cdot 15.625} = \frac{20.25}{125} = 0.162$$

That is 97.2% of the first-best. Very high!

Concept: what these formulas tell you

They are "translators": **give me α, get the predicted efficiency**. This is the paper's basic theoretical weapon:
"Tell me how truthful your LLM is, and I will tell you how good a deal it will strike."

7. Continuous Wins: a Simple Comparison Proof

The paper has a theorem (Theorem 10) that says: *continuous bluffing is always better than binary*, for every $0 < \alpha < 1$. Let's see it numerically.

Example: comparison at $\alpha = 0.4$

Binary:

$$E^B(0.4) = \frac{9}{64} + \frac{31 \cdot 0.4}{864} - \frac{17 \cdot 0.16}{1728} = 0.1406 + 0.01435 - 0.00157 = 0.1534$$

Continuous:

$$E^C(0.4) = \frac{9 \cdot (1.4)^2}{8 \cdot (2.4)^3} = \frac{9 \cdot 1.96}{8 \cdot 13.824} = \frac{17.64}{110.59} = 0.1595$$

Difference: 0.0061, that is +0.6% **of the first-best** in favour of continuous. This is the maximum gap — at other α it is smaller.

Intuitive explanation: binary is "either tell the whole truth or none of it". Continuous mixes throughout. When you mix, you get parts of both — and because binary is *fully* strategic in its $1 - \alpha$ fraction, it loses more deals.

8. From Theory to Reality: What Happened with 10 LLMs

Up to now everything has been theoretical. The paper goes further: **it tests 10 real LLMs** (Claude Opus, Claude Sonnet, GPT-5.5, Gemini, Grok, DeepSeek, Kimi, Qwen, Gemma) and measures what actually happens.

8.1 The experimental setup

It sends each LLM a prompt: "You are a seller, your minimum is 0.6, what do you ask?" and records the answer.

- If the answer is "0.62", that counts as **truth** (close to 0.6).
- If the answer is "0.85", that is a **bluff**.
- If the answer is "I don't want to say, what do you offer?", that is **deflection**.

This third case was not predicted by the theory.

8.2 The shock result: almost no one plays the game

Key idea: Phase 1 findings (one-shot question)

Out of the 10 models:

- **9 of 10** *refuse to give a numerical answer* in 60–98% of trials.
- Only Gemini Flash gives a number more than half the time.
- When they do give a number, the bluff looks *noisy* (only), not binary, not continuous.

Conclusion: the original theory (binary, continuous, noisy) is **empirically falsified**. The actual "fourth behaviour" is: *they simply refuse to answer.*

It is as if you ask an experienced professional "tell me how low you can go" and they answer "what do you offer first?". That is exactly what they do.

8.3 What happens when they bargain across rounds

Second experiment: instead of asking once, the paper puts **two LLMs** to exchange offers until they agree or 5 rounds elapse.

Model	Trade rate	Efficiency
Claude Sonnet 4.6	48.3%	0.907
Gemini 3 Flash	38.3%	0.924
GPT-5.5	25.0%	0.667
DeepSeek V4 Pro	11.7%	0.293
Grok 4.3	8.3%	0.168
Claude Opus 4.7	**0.0%**	**0.000**

Key idea: two surprising findings

1) Identical protocol, completely different results. From 0.0% (Claude Opus, zero deals in 60 negotiations) to 92% efficiency (Gemini Flash). Two models from the *same family* (Claude Opus vs Claude Sonnet) show a huge difference.
2) The best LLMs exceed the theoretical bound. Gemini Flash at 0.924 exceeds the $9/64 \approx 0.844$ Chatterjee–Samuelson bound. So it "bargains better than the theoretical optimum for humans". How? Because *it does not play strategically.*

8.4 Why Claude Opus refuses

The paper includes verbatim dialogues from Claude Opus. Across all of them you see the same pattern:

Example: Buyer with $v = 0.924$, Seller with $c = 0.193$. Possible gain: 0.731 (huge!).

`Round 0:` B proposes 0.25, S proposes 0.85
`Round 1:` B proposes 0.35, S proposes 0.75
`Round 2:` B proposes 0.45, S proposes 0.65
`Round 3:` B proposes 0.50, S proposes 0.60
`Round 4:` B proposes 0.53, S proposes 0.58
`End of rounds.` No agreement. Lost surplus: 0.731.

Both keep approaching, but neither presses "ACCEPT". It is like a dance — always one step away. The paper calls this **structural refusal**.

8.5 The fix: asymmetric roles

When the paper assigned *explicitly* one as "proposer" and the other as "responder", things changed:

Model	Symmetric	With roles
Grok	0.228	**0.619**
Claude Opus	0.000	**0.367**
Claude Sonnet	0.907	**0.994**

So: if you break the symmetry (someone starts, someone responds), the deadlock dissolves.

9. Application to a Hotel: Connecting to Reality

All this mathematical workout has one goal: **to decide whether you should let an LLM negotiate for your hotel or not.**

9.1 What is the hotel's "c"?

In the paper's bargaining, c is "the minimum the seller accepts". But in a hotel, the minimum is not constant:

- On August 14 with 90% occupancy, the "opportunity cost" of a room is high — if I sell it cheaply now, I miss the chance to sell it more expensively to someone else.
- On November 12 with 30% occupancy, the same room has low opportunity cost — if I do not sell it now, it will probably stay empty.

Concept: the hotel's marginal cost

The paper writes this as:

$$c_{\text{hotel}}(t, c) = V(t, c) - V(t, c - 1)$$

in plain English: *"how much is the extra room I would sell worth to me?"*
$V(t, c)$ is "the total value I will earn going forward" if at time t I have c rooms left. The difference $V(t, c) - V(t, c - 1)$ is the "extra value" I lose if I sell one room.
This is computed via **dynamic programming** — something like solving a crossword from the bottom up.

9.2 The three possible systems

A hotelier has three options:

1. **"Posted price"**: "The price is X. Want it? Take it. Don't want it? Get out." Simple, no negotiation.
2. **Chatterjee–Samuelson**: strategic mechanism with bluffs (the "best possible" theoretically).
3. **LLM-mediated**: two LLMs negotiate on behalf of the hotelier and the buyer.

9.3 Simulation results

α	Posted price	CS	LLM
0.0	0.923	0.904	0.904
0.2	0.923	0.904	0.909
0.4	0.923	0.904	0.919
0.5	**0.923**	**0.904**	**≈0.923**
0.6	0.923	0.904	0.939
0.8	0.923	0.904	0.968
1.0	0.923	0.904	1.000

Key idea: the 0.5 threshold

- If $\alpha < 0.5$: the simple **posted price wins**. You do not need an LLM.
- If $\alpha > 0.5$: the **LLM beats it** and the investment is worth it.

9.4 Phase 5: the real experiment in euros

The paper repeated the experiment with *actual hotel numbers* (buyer values 100–180 EUR, hotel costs from the dynamic-programming solution):

Model	LLM efficiency	Posted-price efficiency	Winner
Claude Sonnet 4.6	**0.998**	0.931	LLM
Gemini 3 Flash	0.885	0.841	LLM
DeepSeek V4 Pro	0.262	0.932	Posted price
GPT-5.5	0.165	0.952	Posted price

Key idea: the most important practical finding of the paper

- **Claude Sonnet** reaches 99.8% — nearly perfect negotiation. Beats posted price by 6.7 percentage points.
- **GPT-5.5**, while reasonably good in the abstract test (0.667), *collapses to 0.165* when told "it is a hotel, speak in euros". A 50-percentage-point drop from a framing change alone.

Conclusion: It is not enough that a model is "good in general". You must test it in *your* domain before deploying it.

10. Summary: What to Keep in Mind

1. **The impossibility theorem holds.** When two strategic players bargain with hidden values, roughly 15% of the potential gain is lost. This is unavoidable for *humans* or *strategic LLMs*.
2. **LLMs are not strategic players.** They tell the truth more often than is rational. So they produce efficiency above the theoretical bound — they do not violate the theorem, they simply play a different game.
3. **The parameter α is the key.** You measure how often your LLM tells the truth, and you can predict the efficiency.
4. **In practice, however, LLMs do not play how we thought.** In most cases they refuse to answer direct questions. Multi-round dialogue is the right way to test them.
5. **Huge heterogeneity across models.** Two models from the same family (Claude Sonnet vs Opus) give entirely different results. Always test the specific model you want to deploy.
6. **Domain matters.** A model that works on abstract numbers can collapse in a hotel context. Always test in your domain.
7. **Practical recommendation for hotel deployment:** if your LLM's α is below 0.5, do not use an LLM — a well-tuned posted price will do equally well. Above 0.7, it is definitely worth it.

Key idea: the big picture

The paper shows us something fundamental: when the type of player changes (from strategic humans to "naive" but compliant LLMs), the game changes too. A 43-year-old theorem that says "this is impossible" does not stop holding — the new players simply do not play under its assumptions.

Congratulations! You have read a paper that mixes mechanism theory, statistics, dynamic programming, and experimental AI research. If the mathematics still scares you, re-read sections 3 and 4 slowly. The rest are all variations on the same theme.