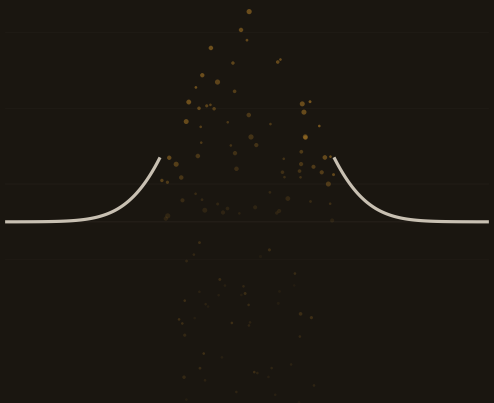


A STATISTICAL SATIRE

THE MEDIAN IS DEAD

A Eulogy for the Reliably Adequate

NABIL EL MAHYAOUI

The Median Is Dead

The Median Is Dead

A Eulogy for the Reliably Adequate

Nabil El Mahyaoui

A statistical satire

Copyright © 2026 Nabil El Mahyaoui.
All rights reserved.

Contents

1 The Bell Curve's Eulogy **1**

1.1 The Architecture of Adequacy 3

1.2 The Comfort of the Middle. 5

1.3 The Day the Median Died. 6

2 The Median Machine **10**

2.1 The Mathematics of Good Enough 11

2.2 The Alignment Tax. 13

2.3 The Temperature Dial 14

2.4 The Price of Being Replaceable. 15

3 The Hollowing **18**

3.1 Who Survives at the Top 19

3.2 Who Survives at the Bottom 21

3.3 The Sixty Percent 23

3.4 The Cruel Precision 25

3.5 The Meritocracy Trap 26

Epilogue: Obituary **29**

References **35**

1

The Bell Curve's Eulogy

You were a solid hire. Not a brilliant one. Your interviewer didn't walk out of the room and call the CTO. Nobody fought over your offer letter. There was no bidding war, no counter-offer from Google, no breathless Slack message in the hiring channel saying "we CANNOT lose this person."

You were a solid hire.

Your resume was clean. Your answers were competent. You whiteboarded a binary search without hesitating and asked two thoughtful questions about the team's deployment pipeline. The hiring committee scored you 3.4 out of 5. The debrief lasted seven minutes.

"Strong yes from everyone. Not spectacular, but exactly what we need."

You started on a Monday. You were productive by the second week. You attended standups, wrote pull requests that passed review on the second try, and remembered to update the Jira tickets. Your manager described you as "dependable." Your peers described you as "solid." At your first annual review, you received the phrase that would define the next fifteen years of your career:

Meets Expectations.

You met them. Every quarter, every cycle, every year. You met expectations the way gravity meets the ground—reliably, invisibly, without anyone pausing to admire the consistency.

You were the median.

And for a while—for decades, actually—the median was the most valuable thing in the economy.

1.1 The Architecture of Adequacy

Let's be precise about what you were, because precision matters when you're writing a eulogy.

The median is not the average. The average is distorted by extremes—by the one engineer who ships ten features a sprint and the one who has been "refactoring" the authentication module since Q2 of last year. The median is the person in the exact middle. Half the team is better. Half the team is worse. And you—you are right there, in the centre of the distribution, doing your job.

You are the developer who writes code that works. Not elegantly—the senior architect will refactor it in six months, tutting gently about abstraction layers. But it works. The tests pass. The feature ships. The customer doesn't notice, which is the highest compliment that adequate code can receive.

You are the lawyer who reviews contracts. Not the partner who structures the deal, not the first-year who misses the clause—the one in the middle who reads every page, flags the standard risks, and bills 1,800 hours

a year, almost all of them doing exactly what you were trained to do.

You are the designer who produces clean, professional layouts. Not the creative director who wins awards, not the intern who can't kern. The one who opens Figma at 9 a.m., follows the design system, and delivers screens that the client describes as "exactly what we asked for."

You are the radiologist who reads sixty scans a day and gets 95% of them right. Not the researcher who publishes breakthrough detection algorithms. Not the one who missed the tumour. The one who does the job, day after day, with a consistency that the entire healthcare system depends on.

The median. The backbone. The sixty percent that holds everything together while the top ten percent gets the keynotes and the bottom ten percent gets the performance improvement plans.

Nobody writes songs about the median. Nobody writes eulogies either.

Until now.

1.2 The Comfort of the Middle

You never understood how good you had it.

The people at the top lived in terror. Every quarter was a referendum on their brilliance. Every project was an existential test. They were only as valuable as their last breakthrough, and the distance between genius and irrelevance was one bad quarter.

The people at the bottom lived in a different kind of terror—the slow, grinding anxiety of performance reviews scored in red, of one-on-ones where the manager's tone shifts to that particular register of rehearsed compassion that precedes a PIP.

You had neither. You had the most dangerous comfort available to a professional in the twenty-first century: the assumption that adequacy would always be enough.

Your annual reviews were pleasant. Your manager ticked the "Meets Expectations" box with the brisk efficiency of someone filling out a form they've filled out a hundred times. You received a cost-of-living adjustment that was technically a raise and functionally a rounding error. You were told you were "valued,"

which in corporate English means "we have no immediate plans to fire you but would not counter-offer if you left."

You stayed. Of course you stayed. The middle is comfortable. The commute is familiar. The health insurance is adequate. The 401(k) match is 4%, which your financial advisor—a median financial advisor, naturally—described as "solid."

You built a life on "solid." A mortgage. Two kids. A car that Consumer Reports rated "above average in reliability," which you noted with quiet satisfaction, because above average is the median's version of ambition.

The foundation was stable. The foundation was, in retrospect, a sheet of ice over warm water. You just couldn't hear the cracking.

1.3 The Day the Median Died

You didn't see it coming, which is the most median thing about you.

The 10x engineers saw it. They'd been playing with the models since the beta, building agents, writing compilers that wrote compilers. They weren't threatened. They were accelerated. The AI was a power tool, and they were already powerful.

The underperformers saw it too, in their own way. They had always survived on institutional inertia—the sheer inconvenience of firing someone and hiring a replacement. They had a different problem, but it wasn't new.

You saw it and assumed it didn't apply to you. You were, after all, *good at your job*. Not so good that you were irreplaceable, but good enough that replacement seemed unnecessary. You occupied that comfortable middle ground where competence is its own justification.

The memo arrived on a Tuesday. Not a dramatic memo—no "we regret to inform you," no "paradigm shift in our operating model." It was, fittingly, a perfectly adequate corporate communication:

> *As part of our ongoing commitment to operational excellence, we are restructuring*

several teams to better leverage AI-assisted workflows. Some roles will be redefined. A small number will be eliminated. Affected employees will be notified individually.

You were notified individually.

The irony—and there is always irony, because the universe has a sense of humour pitched exactly at the frequency of human suffering—is that the memo was indistinguishable from something a competent HR director would write. It was competent, clear, and emotionally calibrated. It expressed regret without liability and optimism without specificity.

It was, in every measurable way, a median piece of corporate prose. The kind of writing that used to require a salaried professional and now requires a prompt.

The AI did not come for the best. The best are safe—for now—because their value lies in the unpredictable, the creative, the strategic leaps that no training data can anticipate.

The AI did not come for the worst. The worst are safe because they are cheap, because their jobs involve phys-

ical presence or human chaos or institutional knowledge so tacit it has never been written down.

The AI came for you. The reliable, competent, meets-expectations, exactly-what-we-need you.

Because you are what the AI *is*.

You are the median. And the machine is a median engine, running at the speed of light, for the cost of electricity.

2

The Median Machine

Here is a fact that the AI industry would prefer you not think about too carefully: large language models are, by construction, machines for producing median output.

This is not a criticism. It is a mathematical description. And like all mathematical descriptions, it is both precise and devastating.

The entire training process—every terabyte of data, every billion-dollar GPU cluster, every carefully tuned reward model—is engineered, from first principles, to produce output that is indistinguishable from what a competent, well-informed, slightly above-average human would produce.

Not what a genius would produce. Not what an idiot would produce. What *you* would produce, on a good day, with unlimited time and perfect recall.

The median. Automated. At scale.

2.1 The Mathematics of Good Enough

Let's open the hood.

A language model is trained by predicting the next token in a sequence. Given all the words that came before, what word comes next? The model sees trillions of examples from the internet—every Wikipedia article, every Stack Overflow answer, every blog post, every legal brief, every radiology report that was ever digitised.

And it learns the most *likely* next word.

Not the most brilliant word. Not the most creative, surprising, or insightful word. The most likely. The one that, across all the millions of humans who have written in similar contexts, appears most frequently in this position.

The training objective is called *cross-entropy loss*, and its mathematical effect is elegant in its cruelty: the model is penalised for deviating from the consensus. If a thousand radiologists would write "no acute findings," the model learns to write "no acute findings." If eight hundred copywriters would open with a rhetorical question, the model opens with a rhetorical question. If the median programmer would use a for-loop, the model uses a for-loop.

Every gradient descent step pushes the model closer to the centre of the distribution. Every parameter update sands down an edge, averages away an outlier, converges toward the mean.

> *Training a large language model is the most expensive averaging operation ever performed.*

Two billion dollars of compute, ten thousand GPUs running for months, the electricity output of a small city—all to produce a machine that can, with superhuman speed and consistency, generate the most statistically ordinary response to any given prompt.

This is not a flaw. It is the *design*.

2.2 The Alignment Tax

It gets worse.

After pre-training—after the model has already absorbed the statistical median of all human text—it goes through a second process called *alignment*. Specifically, Reinforcement Learning from Human Feedback, or RLHF.

Here's how it works. The model generates multiple responses to the same prompt. Human raters rank them. A reward model learns to predict which responses humans prefer. The language model is then fine-tuned to produce responses that maximise the reward.

And what do human raters prefer?

Not brilliance. Brilliance is polarising. A daring, unconventional answer might score a 9 from one rater and a 3 from another. The training process averages across raters. Polarising outputs are penalised.

Not mediocrity. Obviously bad answers score low across the board.

What survives is the response that most people would rate 7 out of 10. Helpful. Accurate. Well-structured.

Appropriately hedged. Fluent without being flashy. Correct without being provocative.

The 7/10 response. The B+. The "good job, but—"

RLHF does not optimise for excellence. It optimises for the *absence of objection*. And the absence of objection is, by definition, the median of human preferences.

The aligned model is a machine that has been trained, at enormous expense, to never be the worst thing in the room and never be the best. To be exactly, precisely, immaculately *adequate*.

Sound familiar?

2.3 The Temperature Dial

There is a parameter called *temperature* that controls the randomness of the model's output.

At temperature 0, the model always picks the single most likely token. Pure, undiluted median. The platonic ideal of "what most people would say."

At temperature 1, the model samples from the full distribution—occasionally brilliant, occasionally in-

coherent, always unpredictable. This is jazz. Nobody ships jazz to production.

At temperature 0.7—the industry default, the setting that ships in virtually every commercial API—the model is creative enough to vary its phrasing and conservative enough to never embarrass anyone. It is the setting that produces output indistinguishable from a competent professional on a Tuesday afternoon.

You can, quite literally, adjust how median the output is with a single number.

The industry settled on 0.7 because focus groups found it "helpful, clear, and professional." The same three words your manager used in your last performance review.

The temperature dial is not a technical parameter. It is a knob that controls how much the machine sounds like you. And 0.7 sounds *exactly* like you.

2.4 The Price of Being Replaceable

Here is the arithmetic that kills the median.

You cost $85,000 per year, plus benefits, plus office space, plus management overhead, plus the two weeks in February when you have the flu and your tickets pile up. Fully loaded, you cost your employer roughly $130,000 per year to produce output that meets expectations.

The model costs $20 per month.

It produces output that also meets expectations. It does not have the flu. It does not have a commute. It does not need a one-on-one with its manager to discuss career development. It does not quietly update its LinkedIn profile during a reorg. It does not require a farewell cake when it leaves, because it does not leave.

The quality gap between your output and the model's output is, on average, zero. Not because the model is brilliant. Because you are—and I say this with the deepest respect for everything you've built—*average*. You are exactly as good as the statistical composite of every person who has ever done your job. And the model *is* that composite.

You are being replaced not by something better than you. You are being replaced by a $20/month version of yourself.

Your manager will describe the transition as "augmentation, not replacement." This is a lie, but it is a well-intentioned lie, and it tested well with the internal comms team. They rated the messaging 7 out of 10.

3
The Hollowing

The economy has a shape, and the shape is changing.

For seventy years—since the GI Bill flooded universities, since the post-war boom created the professional class, since "go to college, get a good job" became the incantation of upward mobility—the shape of the economy was a bell curve. A few at the top. A few at the bottom. And a vast, comfortable, mortgage-paying, vacation-taking, pension-expecting middle.

That middle was you.

The new shape is a barbell.

Two weights at the extremes, connected by a thin bar where the middle used to be. On one side: the irre-

placeable. On the other: the un-automatable. In between: a void.

You are falling into the void.

3.1 Who Survives at the Top

The top five percent was never in danger, and it's important to understand why, because the reason is not what LinkedIn influencers tell you.

It is not "creativity." The model is creative. It generates poetry, composes music, designs interfaces, and writes fiction that is—by every empirical measure—indistinguishable from competent human output. Telling a displaced mid-level professional to "be more creative" is like telling someone drowning in the Atlantic to "be more buoyant." It is advice that mistakes the vocabulary of survival for a strategy.

The top survives because its value is generated at the *frontier* of the domain—in the space where no training data exists, where the next move cannot be predicted from the last ten thousand moves.

The architect who looks at a system and sees the redesign that nobody proposed survives because the redesign has no precedent for the model to learn from. The trial lawyer who reads a jury and shifts strategy mid-cross-examination survives because juries are chaotic systems and no corpus of transcripts captures the dynamics of twelve confused strangers deciding someone's fate. The surgeon who encounters the unexpected anatomy survives because the unexpected, by definition, is not in the distribution. The founder who sees a market that doesn't exist yet survives because you cannot predict the next token of a sentence that hasn't been started.

The top five percent operates in the tails of the distribution—the region where the bell curve approaches zero, where the model's confidence collapses, where the median has nothing to offer.

They are safe. But their safety is not transferable. You cannot become a top-five-percent performer by attending a weekend workshop on "AI-proof your career." The top is the top because it cannot be taught. It can only be discovered, through years of deliberate practice, through failures that the median never risked,

through the accumulation of judgement so deep it looks like intuition.

If you are reading this and thinking "I am in the top five percent," I would gently suggest that the top five percent is too busy doing the work to read an essay about whether the work will survive.

3.2 Who Survives at the Bottom

The bottom twenty percent is safe for an entirely different reason: economics.

A plumber survives not because AI cannot understand plumbing. It can. Ask any frontier model to diagnose a pressure drop in a residential hot-water loop and it will give you a better differential diagnosis than most apprentices. The plumber survives because the AI cannot crawl under your house with a wrench.

A home health aide survives not because emotional care is beyond computation—the model is arguably more patient and empathetic than most humans on a twelve-hour shift. The aide survives because the patient needs a body in the room. A physical presence that can catch someone when they fall, that can hold

a hand, that can exist in three-dimensional space with mass and warmth and the faint smell of coffee.

A line cook at a diner survives because the diner serves 400 meals a day at $8.50 each, and the robotics required to replace the cook cost more than the cook does. The economics don't work. Not because the technology doesn't exist, but because the margin doesn't justify the capital expenditure.

The bottom survives on embodiment, on physicality, on the stubborn fact that atoms are harder to manipulate than bits. The AI can write the recipe. It cannot chop the onion.

This is not a dignified survival. It is an economic reprieve. And it lasts only as long as the cost of robotics exceeds the cost of human labour—a gap that is narrowing, year by year, but slowly enough that the bottom has time.

Time to do what, exactly, is a question nobody has answered.

3.3 The Sixty Percent

And then there is you.

You are not at the top. Your work, however competent, is not at the frontier. It is reproducible. It follows patterns. It is, in the precise technical sense, *within the training distribution*.

You are not at the bottom. Your work is not physical. You do not crawl under houses or hold hands or chop onions. You sit at a desk and produce information-based output—reports, code, designs, analyses, recommendations—that can be transmitted as bits.

You are in the middle. The wide, comfortable, doomed middle.

The associate who reviewed contracts for two hundred billable hours last month. The AI reviewed the same documents in eleven minutes. The partner checked both outputs. Yours had three missed clauses. The AI had none. You had coffee stains on page 47.

The designer whose Figma file has 340 frames and a naming convention only you understand. The AI's file

has 340 frames and a naming convention that everyone understands. You spent three days. The AI spent four minutes. Your creative director called both versions "fine," which is the word people use shortly before eliminating your position.

The content writer who reliably produced 800 words on any topic, in any tone, by end of day. The AI produces 800 words on any topic, in any tone, in four seconds. Your editor cannot tell the difference. Neither can you, which you discovered when you were asked to identify which blog posts were yours and got three out of ten wrong.

The analyst who built dashboards. The project manager who updated the Gantt chart. The consultant who wrote the slide deck. The accountant who filed the returns.

These are not bad jobs. These are not lazy people. They are competent professionals doing exactly what they were trained to do. And that is the problem. They were trained to produce output that falls within the distribution. And the distribution is exactly what the model has learned.

3.4 The Cruel Precision

Here is what makes this different from every previous wave of automation.

The Industrial Revolution replaced muscle. It targeted the physical middle—the weaver, the spinner, the carder. But it created a new middle: the machine operator, the factory foreman, the clerk. The knowledge economy absorbed the displaced because the new machines needed new humans to run them.

The AI revolution replaces *judgement*. Not physical judgement—the surgeon's hand, the pilot's reflexes. Informational judgement. The ability to read a document and decide what matters. The ability to look at data and identify the pattern. The ability to take a requirement and produce a competent solution.

This is exactly what the middle does. This is its entire value proposition.

And unlike the Industrial Revolution, the new machines do not need new humans. The AI does not require operators. It does not need foremen. It scales with electricity, not headcount. Every improvement in the model makes it *more* autonomous, not less. The

trend line points toward fewer humans in the loop, not more.

The automation of muscle created the knowledge economy. The automation of median judgement creates—what, exactly?

Nobody has a good answer. And the absence of that answer is the sound of sixty percent of the professional class holding its breath.

3.5 The Meritocracy Trap

The cruelest part is that you did everything right.

You went to college. You got the degree. You learned the skills. You showed up, every day, and did your job competently. You followed the script that society handed you at age eighteen and executed it faithfully for twenty years.

The script said: competence is rewarded. The script said: if you are reliable, you are valuable. The script said: the middle is safe.

The script lied.

Not deliberately. When the script was written—in the 1960s, in the 1980s, in the 2000s—it was true. Competence *was* rewarded, because competence was scarce enough to command a price. A person who could write clear prose, or review a contract, or build a database, or design a landing page was worth $85,000 a year because the supply of such people was limited by the time it took to train them and the biological constraints of a single human working eight hours a day.

The model has no such constraints.

It was trained once, at great expense, and now it produces median-quality output for every skill simultaneously, in every language, at every hour, at a marginal cost approaching zero. The supply of competence just went from scarce to infinite. And when supply becomes infinite, price goes to zero.

This is not a temporary disruption. It is not a recession you can wait out. It is the economic equivalent of a phase transition—a permanent change in the state of the labour market, as irreversible as ice melting into water.

The bell curve is not shrinking. It is being *hollowed*. The tails remain. The middle dissolves.

And the sound it makes as it dissolves is the sound of a million keyboards going quiet in open-plan offices around the world, replaced by the soft, efficient hum of inference servers doing the same work, at the same quality, for a fraction of a cent.

Epilogue

THE PROFESSIONAL REGISTER

OBITUARY

THE MEDIAN PROFESSIONAL

b. 1952 — d. 2026

The Median Professional died quietly on a Tuesday morning, at a desk, in an open-plan office, during a company-wide restructuring that was described by management as "an exciting evolution of our operating model."

The cause of death was officially recorded as "redundancy," though the attending algorithms noted that a more precise diagnosis would be "statistical replaceability at negative return on investment."

The Median Professional was 74 years old. Colleagues remarked that the role had been "showing its age" since at least 2023, though no one had thought to say so while the invoices were still being paid.

Born in the post-war expansion to Institutional Meritocracy and the GI Bill, the Median Professional was raised on a simple promise: that competence, reliably delivered, would be exchanged for stability, reliably provided. This promise was honoured for approximately sixty years, which in retrospect was less a social contract and more of a promotional offer with an undisclosed expiration date.

The Median Professional attended a respectable university—not the best, not the worst—and graduated with a degree that was, at the time, considered sufficient. Early career achievements included: learning the filing system, mastering the photocopier, and developing an intuitive grasp of which emails required a reply and which could be safely ignored. These skills, once considered essential, are now performed by a cloud-based API that costs less per month than the Median Professional's daily coffee.

Over a career spanning four decades and four employers, the Median Professional held a series of titles—Analyst, Associate, Manager, Senior Manager—that conveyed a sense of upward momentum while describing an essentially lateral trajectory.

Performance reviews consistently noted that the Median Professional "met expectations," "delivered solid work," and "was a reliable team player." At no point did any review use the word "exceptional," which, as it turned out, was the only word that would have mattered.

The Median Professional's greatest accomplishment, by all accounts, was consistency. In thirty-eight years of employment, the Median Professional never once produced work that was complained about, celebrated, or remembered. This record of uninterrupted adequacy was, until very recently, considered a virtue.

The illness was first detected in late 2022, when a chatbot passed a bar exam that the Median Professional's nephew had failed twice. Symptoms progressed rapidly. By early 2024, the Median Professional's output was routinely indistinguishable from

that of a $20-per-month subscription service—a comparison that was made explicitly, in a boardroom, by a vice president who had been the Median Professional's mentee.

The Median Professional sought treatment in the form of an online course titled "Future-Proof Your Career with AI." The course was twelve hours long, cost $499, and consisted almost entirely of platitudes about "embracing change" and "leveraging your uniquely human skills," none of which were specified. The Median Professional completed it and received a certificate of completion, which was added to a LinkedIn profile that continues to receive automated endorsements for "Microsoft Excel" and "Team Player."

The certificate did not help.

The Median Professional is survived by:

- **The Top Five Percent,** who expressed sincere condolences while quietly absorbing the Median Professional's responsibilities at 1.5 times the salary, assisted by the same tools that performed the displacement.

- **The Bottom Twenty Percent,** who were briefly alarmed before being reassured that their roles "require a physical presence that AI cannot replicate at this time." The qualifier "at this time" was noted but not discussed.
- **A 401(k)**, now valued at slightly less than what financial advisors describe as "enough," and slightly more than what dignity requires.
- **Two children,** both of whom are pursuing careers in fields that did not exist when the Median Professional started working, and may not exist by the time they finish their degrees.
- **A LinkedIn profile** that remains active, continues to auto-generate content, and has gained 340 followers since the Median Professional stopped posting personally, a statistic that the family has asked not to be mentioned but which is included here because it is, in its own quiet way, the most devastating detail of this obituary.

The Median Professional was predeceased by the Gold Watch (d. 1998), the Defined-Benefit Pension (d. 2008), the Concept of Job Security (d. 2020), and

the Assumption That a Degree Guarantees Employment (d. 2024), all of whom passed within the same generation under broadly similar circumstances.

A memorial service will be held on the company's internal Slack channel, in the `#general` room, between the Q3 earnings announcement and a reminder about the updated PTO policy. Attendance is optional but will be noted by the analytics dashboard.

In lieu of flowers, the family requests that you upskill.

References & Further Reading

Autor, D. (2024). "Applying AI to Rebuild Middle Class Jobs." NBER Working Paper No. 32140.
David Autor has spent two decades studying how technology reshapes labour markets. This paper argues AI could rebuild the middle class. We admire his optimism the way one admires a sandcastle built at high tide.

Kaplan, J. et al. (2020). "Scaling Laws for Neural Language Models." *arXiv:2001.08361*.
The paper that quantified how model capability scales with compute and data. It does not mention what scales inversely: the economic value of being competent.

Graeber, D. (2018). *Bullshit Jobs: A Theory*. Simon & Schuster.
Graeber argued that many professional roles serve no

meaningful purpose. AI agreed, and has begun acting on the finding.

Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable*. Random House.
On the fragility of systems that mistake the absence of catastrophe for the presence of safety. The median professional's career was a Thanksgiving turkey—well-fed every day, right up until the day that mattered.

Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age*. W. W. Norton.
Predicted that digital technologies would hollow out the middle of the labour market. They were right. They were also a decade early, which in the technology prediction business counts as nearly correct.

Ouyang, L. et al. (2022). "Training Language Models to Follow Instructions with Human Feedback." *NeurIPS 2022*.
The RLHF paper. It describes, in meticulous technical detail, how to train a machine to produce responses that the median human rater considers "helpful." The paper does not explore what happens to the median human.

Markovits, D. (2019). *The Meritocracy Trap*. Penguin Press.
On how meritocracy concentrates advantage at the top and anxiety everywhere else. Markovits argues that the system is rigged. AI is not rigging it differently. It is simply making the rigging faster.

About the Author

Nabil El Mahyaoui

https://nabil-el-mahyaoui.com/

https://www.linkedin.com/in/elmahyaoui