RAFAEL MORAES BATISTA

Email: rbatista@princeton.edu
Website: rafaelmbatista.com
Bluesky: rafmbatista.bsky.social

Carey Business School
Johns Hopkins University
325 7th St NW Washington, DC 20004

RESEARCH TOPIC AREAS

Psychology & Algorithms; Scientific Discovery; Financial Decision-Making; Judgment & Decision-Making; Consumer Language; Lab and Field Experiments

ACADEMIC POSITIONS

Postdoctoral Research Fellow, **Johns Hopkins Carey Business School** — Washington, DC
Technology & Society Initiative — 2026–

Postdoctoral Research Associate, **Princeton University** — Princeton, NJ
Kahneman-Treisman Center for Behavioral Science & Public Policy — 2025-2026

EDUCATION

The University of Chicago Booth School of Business — Chicago, IL
PhD, Behavioral Science — 2019-2025
Dissertation Committee: Sendhil Mullainathan (Co-Chair), Abigail Sussman (Co-Chair), Alex Imas, Jean-Pierre Dube

The University of Chicago Booth School of Business — Chicago, IL
MBA — 2023-2025

Princeton University — Princeton, NJ
Visiting Student, Psychology and Princeton School of Public & International Affairs — 2022
Advisor: Elizabeth "Betsy" Levy Paluck

University of Warwick — Coventry, UK
MS, Behavioural & Economic Science (Distinction) — 2014-2015
Advisors: Daniel Read & Nick Chater

Florida State University — Tallahassee, FL
BS, Psychology (Honors) — 2010-2014
Advisor: Roy Baumeister
BA, International Affairs (concentration in Economics)
Advisors: John Hamman & Darrin McMahon

Koç University, Istanbul — Istanbul, Turkey
Visiting Student — 2014
Advisor: Çigdem Kagitçibasi

JULY 2026
Latest CV available at: rafaelmbatista.com/cv

PROFESSIONAL EXPERIENCE

Independent Consultant, UNICEF 2026

Vendor, Barraca da Ana 104 (Field Work, Ipanema Beach) 2025

PhD Fellow, UN Behavioral Science Team 2022

Research Professional, University of Chicago Booth School of Business 2019

Head of Behavioral Research, Commonwealth Bank of Australia 2016-2019
Senior Behavioral Scientist, Commonwealth Bank of Australia

- Yu, J., Sanskriti Merchant, Bing Feng and Dilip Soman. "Banking on behavioral science: Commonwealth Bank of Australia." Behavioral Economics in Action at Rotman Case Study, October 2023.
- Buell, Ryan W., and Leslie K. John. "Commonwealth Bank of Australia: Unbanklike Experimentation." Harvard Business School Case 619-018, October 2018. (Revised February 2020.)

Research Associate, Busara Center for Behavioral Economics 2015-2016

Research Assistant, Baumeister & Tice Social Psychology Lab 2012-2014

PUBLICATIONS AND WORKING PAPERS

Batista, Rafael M., Abigail B. Sussman, and Jennifer Trueblood (2023) "Self-other differences in perceptions of wealth", *Journal of Experimental Social Psychology*, https://doi.org/10.1016/j.jesp.2022.104420

Batista, Rafael M. and James Ross (2024) "Words that work: Using large language models to generate and refine hypotheses," *Invited Revision at Journal of Consumer Research (JCR),* https://dx.doi.org/10.2139/ssrn.4926398

- Winner of *Psychology of Tech Dissertation Awards 2026*

Batista, Rafael M., Ella Mao, and Abigail B. Sussman (2025). "Keeping cash and revolving debt: How consumers' preference for spending on debit versus credit influences their decision to co-hold," *Invited revision (2nd Round) at Journal of Marketing Research (JMR),* https://dx.doi.org/10.2139/ssrn.4558490

Batista, Rafael M., Ella Mao, Abigail B. Sussman, Neale Mahoney and Jessica Min (2025). "Disclosing the costs of co-holding liquid assets and high-interest debt has limited impact on behavior," *Invited revision (2nd Round) at PNAS Nexus,* http://dx.doi.org/10.2139/ssrn.5364837

Batista, Rafael M., Juliana Schroeder, Aastha Mittal, and Sendhil Mullainathan (2024) "Misarticulation: Why we sometimes feel our words don't match our thoughts," *Working paper*, https://dx.doi.org/10.2139/ssrn.4687986

Batista, Rafael M. and Tom Griffiths (2026) "A Rational Analysis of the Effects of Sycophantic AI," *Working paper,* https://doi.org/10.48550/arXiv.2602.14270

WORK IN PROGRESS

Norm Querying (with Betsy Levy Paluck)

- *Descriptive survey data collected*

Social Norms Field Study on Reducing Social Media Use Among Teenagers (with Leonardo Bursztyn, Betsy Levy Paluck, Cindy Ha, and The Behavioural Insights Team)
- *53 schools confirmed*
- *Baseline conducted: Nov 2025*
- *Intervention released: Nov 2025*
- *Endline conducted: Feb 2026*

Tamos Juntos: Marketing Strategy in Informal Markets
- *>100 hours of intensive ethnography and participant observation*

A Psychology of Stasis: A Theory of Behavioral Maintenance
- *Early-stage; Current draft focuses on theoretical framework with studies planned*
- *Grant received (with Yena Kim) for $4,500 to conduct studies*

Reasons versus Rationalizations: Untangling justifications for everyday decisions (with Jane Risen)
- *Several studies conducted*

PROPRIETARY RESEARCH

Choosing chickens (Researcher; Ethiopia). A mixed-method research project in the field to understand how rural farmers chose chickens in rural market settings.

Unconditional cash transfers (Project Manager; Kenya). A randomized controlled trial to assess the effects of unconditional cash transfers to low income households in rural Kenya. See Haushofer and Shapiro (2016)

Decreasing loan default rates (Researcher; Kenya). An incentive-compatible artefactual field experiment which simulated borrowing to test various interventions to decrease default rates.

Increasing savings for health emergencies (Researcher; Kenya). A series of pilots designed to test interventions to increase savings earmarked for health emergencies amongst low income households.

Measuring willingness-to-pay for insurance (Researcher; Kenya). A series of framed field experiments to generate valuations of innovative insurance products in emerging markets.

Intrahousehold bargaining power (Consultant; Kenya). A series of framed field experiments to measure spousal bargaining power in informal urban settlements in Nairobi and rural Kenya.

Improving hygiene and reducing the use of defecation (Consultant; Kenya). A behaviourally-informed communications campaign to increase use of clean toilets throughout Kenya's informal settlements.

Refunds and bonuses: A tax time replication study (Researcher; Australia). A randomized controlled trial with 500,000 customers testing the effects of framing on savings decisions.

Transparency of trade-offs to improve credit card selection online (Consultant; Australia). A randomized controlled trial with 400,000 customers testing the effects of presenting products' benefits and drawbacks.

Operational transparency for online personal loan applications (Consultant; Australia). A randomized controlled trial assessing the effects of presenting customers with transparency into the servicing of loan applications.

Choosing chickens (Researcher; Ethiopia). A mixed-method research project in the field to understand how rural farmers chose chickens in rural market settings.

Unconditional cash transfers (Project Manager; Kenya). A randomized controlled trial to assess the effects of unconditional cash transfers to low income households in rural Kenya. See Haushofer and Shapiro (2016)

Decreasing loan default rates (Researcher; Kenya). An incentive-compatible artefactual field experiment which simulated borrowing to test various interventions to decrease default rates.

Increasing savings for health emergencies (Researcher; Kenya). A series of pilots designed to test interventions to increase savings earmarked for health emergencies amongst low income households.

Aligning payments to income (Consultant; Australia). A natural experiment examining the impact on financial decisions when large expenses are aligned to income.

Socially responsible algorithms (Consultant; Australia). Using machine learning to develop algorithms that are equitable and commercially viable. See Coston, Rambachan, and Chouldechova (2021)

INVITED TALKS

Presentation on “A Rational Analysis of Sycophantic AI” at the UK AI Security Institute (AISI, virtual)	April 2026
Presentation on “Sycophantic AI and Beliefs” at MIT Human Cooperation Lab (virtual)	February 2026
Presentation on “Words that Work: Combining AI and Behavioral Science to Generate and Refine Hypotheses from Text” at the Technology and Society Initiative, Johns Hopkins University Carey School of Business	January 2026
Presentation on “Words that Work: Combining AI and Behavioral Science to Generate and Refine Hypotheses from Text” at the Marketing Seminar Series, Stanford Graduate School of Business	November 2025
Presentation on “Using language to generate hypotheses” at the Behavioral Decision-Making Seminar Series, UCLA Anderson School of Business	February 2025
Presentation on “Using language to generate hypotheses” at Work and Organizational Studies, MIT Sloan School of Management	January 2025
Presentation on “Using language to generate hypotheses” at Marketing Department Seminar Series, Yale School of Management	November 2024
Presentation on “Using language to generate hypotheses” at the Applied Quantitative Methods Workshop, UChicago Political Science Department	November 2024
Presentation on “Using language to generate hypotheses” at Marketing Department Seminar Series, MIT Sloan School of Management	October 2024
Presentation on “Using language to generate hypotheses” at Marketing Department Seminar Series, Indian School of Business	October 2024
Presentation on “Using language to generate hypotheses” at the Behavioral	October 2024

Decision-Making Seminar Series, UCLA Anderson School of Business

Presentation on "Using language to generate hypotheses" at Marketing Department Seminar Series, UC San Diego–Rady School of Business — October 2024

Presentation on "Using language to generate hypotheses" at Marketing Department Seminar Series, Harvard Business School — September 2024

Presentation on "Using language to generate hypotheses" at the Psychology & Economics Lunch Seminar Series, UC Berkeley Haas School of Business — September 2024

CONFERENCES (underlined to indicate I presented)

Batista, R. M. & Ross, J. *Generating, Refining and Validating Hypotheses: A Systematic Process for Discovering Language Effects on Decisions*. Presentation at SJDM, Denver, CO, USA. — November 2025

Batista, R. M. & Ross, J. *Words that Work: Using Language to Generate and Refine Hypotheses*. Presentation at ACR, Washington, DC, USA. — October 2025

Batista, R. M. & Ross, J. *Words that Work: Using Language to Generate Hypotheses*. Presentation and poster at Human-AI Complementarity for Decision Making, Pittsburgh, PA, USA. — September 2025

Batista, R. M. & Ross, J. *Words that Work: Using Language to Generate Hypotheses*. Presentation at SCP, Las Vegas, NV, USA. — March 2025

Batista, R. M. & Ross, J. *Words that Work: Using Language to Generate Hypotheses*. Presentation at NeurIPS, Vancouver, Canada. https://openreview.net/forum?id=pvnXlajGBZ — December 2024

Batista, R. M. & Ross, J. *Words that Work: Using Language to Generate Hypotheses*. Presentation at CoDE@MIT, Boston, MA, USA. — October 2024

Batista, R. M. & Ross, J. *Words that Work: Combining Machine Learning and Psychology to Generate Hypotheses from Tex*t. Presentation at AOM, Chicago, IL, USA. — August 2024

Li, S., Batista, R. M., & Schroeder, J. *Who's at Fault? Assignment of Responsibility in Miscommunication*. Presentation at AOM, Chicago, IL, USA. — August 2024

Batista, R. M., Mao, E., Sussman, A. *A Psychological Account of Co-Holding: Why Some Choose to Simultaneously Save and Borrow.* Presentation at the Boulder Summer Conference on Consumer Decision-Making, Boulder, CO, USA. — March 2023

Batista, R. M., Mao, E., Min, J., & Sussman, A. *A Psychological Account of Co-Holding: Why Some Choose to Simultaneously Save and Borrow.* Presentation at SCP, San Juan, Puerto Rico. — March 2023

Batista, R. M., Mao, E., Min, J., & Sussman, A. *Understanding Why Customers Co-hold Savings and Debt: A Field Experiment.* Presentation at SJDM, San Diego, CA, USA. — November 2022

Batista, R. M., Mao, E., Min, J., & Sussman, A. *Understanding Why Customers Co-hold Savings and Debt: A Field Experiment.* Presentation at ACR, Denver, CO, USA. — October 2022

Batista, R. M., Mao, E., Min, J., & Sussman, A. *Understanding Why Customers Co-hold Savings and Debt: A Field Experiment.* Presentation at CoDE @ MIT, Cambridge, MA, USA. — October 2022

Batista, R. M., Mittal, A., Schroeder, J., & Mullainathan, S. *Misarticulation: Theory and Evidence.* Presentation at SESP, Philadelphia, PA, USA. — October 2022

Batista, R. M., Mao, E., Min, J., & Sussman, A. *Understanding Why Customers Co-hold Savings and Debt: A Field Experiment.* Presentation at the ISMS Marketing Science Conference, Virtual. — June 2022

Batista, R. M., Mittal, A., Schroeder, J., & Mullainathan, S. *Misarticulation: Theory and Evidence.* Poster presented at Society for Judgment and Decision-Making, Virtual. — February 2022

Batista, R. M., Sussman, A., & Trueblood, J. "Self-Other Differences in Perceptions of Wealth", in *NA - Advances in Consumer Research* Volume 49, eds. Tonya Williams Bradford, Anat Keinan, and Matthew Matthew Thomson, Duluth, MN: Association for Consumer Research, Pages: 370-370. — October 2021

Batista, R. M., Vonasch, A. J., Aldanmaz, B., Baumeister, R. F., & Kagitçibasi, C. *Meanings of Life Across Cultures: Five Needs for a Meaningful Life.* Poster presented at the International Convention of Psychological Science, Amsterdam, The Netherlands. — March 2015

Batista, R. M., Gladstone, J., Nikitin, E., & Ravichandar, N. *When saving, who needs targets anyway?* Poster accepted for the 14th TIBER Symposium on Psychology and Economics, Tilburg, The Netherlands. — August 2015

ADDITIONAL TALKS & PRESENTATIONS

Presentation on "Words that Work: Combining AI and Behavioral Science to Generate and Refine Hypotheses from Text" in LLM & Behavioral Science PhD Seminar (taught by David Rand), Cornell University - Virtual — November 2025

Presentation on "Using language to generate hypotheses" at the Salganik Lab (Matt Salganik), Princeton University — October 2025

Presentation on "Using language to generate hypotheses" at the UChicago Political Science Quantitative Methods Workshop, University of Chicago — November 2024

Presentation on "Using language to generate hypotheses" at the Texas Behavioral Science and Policy Institute (David Yeager), UT Austin — October 2024

Presentation on "Using language to generate hypotheses" at the Knowledge Lab (James Evans), University of Chicago — September 2024

Presentation on "Using language to generate hypotheses" at the Judgment & Decision-Making Lab, UC Berkeley Haas School of Business — September 2024

Presentation on "Generating Hypotheses from Text" at the Center for Decision Research Brown Bag, University of Chicago Booth School of Business — April 2024

Presentation on "Keeping Cash and Revolving Debt" at the Center for Decision Research — October 2023

Brown Bag, University of Chicago Booth School of Business

Presentation on "Misarticulation: Theory and Evidence" at The NERD Lab, Harvard Business School — November 2021

Presentation on "Misarticulation: Theory and Evidence" at the Center for Decision Research Brown Bag, University of Chicago Booth School of Business — October 2021

Presentation on "Misarticulation: Theory and Evidence" at The Schroeder / Critcher Lab, University of California-Berkeley Haas School of Business — October 2021

Presentation on "Misarticulation: Theory and Evidence" at The Social Systems Lab (Thalia Wheatley), Dartmouth — July 2020

Presentation on "Behavioural, Science Not Included," at the Sydney Behavioural Economics Network Meetup, Sydney, Australia — August 2017

Presentation on "Identifying Behavioral Barriers and Interventions for Digitized Services," at the Behavioural Research for Emerging and Applied Knowledge, Nairobi, Kenya — July 2016

TEACHING & TRAINING

Princeton University

Psychology for Policy Analysis and Implementation *(with Betsy Levy Paluck and Anuj Shah, MPA level)* — Spring 2026

University of Chicago Booth School of Business

Managerial Psychology *(with Ann McGill, Executive MBA level)* — Fall 2023

Psychological Tools of Thought *(with Sendhil Mullainathan, PhD level)* — Spring 2023

Artificial Intelligence *(with Sendhil Mullainathan, MBA level)* — Winter 2023, Winter 2022, Winter 2019

Managerial Psychology *(with Ann McGill, Executive MBA level)* — Fall 2022

Behavioral Economics *(with Alex Imas, Undergraduate level)* — Spring 2022

Data Science, Algorithms, and Society *(with Sendhil Mullainathan and Devin Pope, PhD level)* — Winter 2022

University of Chicago Harris School of Public Policy

Artificial Intelligence *(with Jens Ludwig, MPP level)* — Winter 2019

Commonwealth Bank of Australia

Executive Training *(4 hours x 20 sessions)* 2017

WORKSHOPS

Invited- UNICEF-Gates Foundation 2025 Behavioural Horizons 2025

Selected- AI Institute for Societal Decision Making Workshop: Human-AI Complementarity for Decision Making (Pittsburgh, PA, USA) 2025

Selected- Machine Learning in Economics Summer Institute (Chicago, IL, USA) 2024

Selected- International Rationality Summer Institute (Aurich, Germany) 2016

Invited- Compassion Week, Decision Research Institute (Eugene, OR, USA) 2016

AFFILIATIONS

Kahnemen-Treisman Center for Behavioral Science & Public Policy (Princeton)
AI Lab (Princeton)
STAR Lab (Harvard)
Association for Consumer Research (ACR)
Society for Consumer Psychology (SCP)
Society for Judgment & Decision-Making (SJDM)

COMPUTATIONAL SKILLS & PROGRAMMING R; Python; Qualtrics

LANGUAGES English *(Native in Speaking, Writing, and Reading);* Portuguese *(Native in Speaking; Competent in Writing and Reading);* Spanish *(Conversational in Speaking; Competent in Writing and Reading)*

HOBBIES Stand-up comedy; volleyball; running; cooking; dancing, trying other hobbies

---

FORMER MENTEES & DIRECT REPORTS

Jessica Min (Princeton) -jessica.min@oecd.org
Sophie Wang (UChicago) - jiaruiwang@uchicago.edu
Grant Harris (UChicago) - gwharris@uchicago.edu
Aveesha Gandhi (UChicago) - aveeshagandhi@uchicago.edu
Nikhil Gajbhiye (Ashoka University) - nikhil.gajbhiye_ug2023@ashoka.edu.in

---

REFERENCES

Sendhil Mullainathan, Professor of Economics and Peter de Florez Professor of EECS at MIT, sendhil@mit.edu

Abigail Sussman, Professor of Marketing and Beatrice Foods Co. Faculty Scholar at Booth School of Business abigail.sussman@chicagobooth.edu

Tom Griffiths, Professor of Psychology and Computer Science at Princeton University, tomg@princeton.edu

Betsy Levy Paluck, Eugene Higgins Professor of Psychology and Public Affairs, Princeton University, epaluck@princeton.edu

Dilip Soman, Canada Research Chair in Behavioural Science and Economics at Rotman School of Management, dilip.soman@rotman.utoronto.ca

Juliana Schroeder, Harold Furst Chair in Management Philosophy and Values Professor at UC Berkeley Haas School of Business, jschroeder@haas.berkeley.edu

---

ABSTRACTS

Batista and Ross (2024, Invited revision at *JCR*)

In this paper, we examine how specific features of language drive consumer behavior. Our contribution, however, lies not in testing specific hypotheses; rather, it is in demonstrating a data-driven process for generating them. We introduce a three-step framework—Hypothesize, Intervene, and Predict—that integrates large-language models (LLMs), machine learning (ML), and experimental validation to systematically discover and prioritize testable insights from text. Using a dataset with over 60,000 headlines (and over 32,000 A/B tests), we first Hypothesize linguistic features by prompting an LLM to identify differences between headline pairs. We then use an LLM to Intervene–systematically rewriting headlines to incorporate these features–and use an ML model to Predict the causal impact of these changes on consumer engagement. The framework generates a prioritized list of hypotheses, which we validate on a hold-out set of 1,600 A/B tests. Our approach indeed facilitates discovery. For instance, we find that describing physical reactions significantly increases engagement. In contrast, focusing on positive aspects of human behavior decreases it. This approach extends beyond a single application. In general, it offers a data-driven method for discovery and refinement that is capable of converting unstructured text data into insights that are interpretable, novel, testable, and generalizable. It does so while maintaining a transparent role for both human researchers and algorithmic processes. This approach offers a practical tool to researchers, organizations, and policymakers seeking to aggregate insights from multiple marketing experiments.

Batista, Sussman, and Trueblood (2023, *JESP*)

People evaluate their own wealth differently from how they evaluate the wealth of others. Across six experiments, we find evidence that people focus disproportionately on debt when thinking about their own (vs. another person's) wealth. In Experiments 1–3, participants predicted how wealthy they or someone else would be in one year, assuming they had the same amount in assets and debt today. While participants were generally optimistic about the future, they believed debt would shrink faster for themselves than for others. Participants focused more on paying down debt than growing assets when thinking about their own wealth. Further, when asked to consider what they would do with a windfall, they allocated more towards repaying debt than they believed others would. In Experiments 4 and 5, participants assessed their own wealth or that of another person after purchasing a car or a house and borrowing to do so. In every case, participants considered others (vs. oneself) as better off financially when holding the price and amount borrowed constant. As debt increased, the gap between self and others widened. In Experiment 6, a separate group of participants also reported their beliefs about how others might see them. When actively considering another person's perspective, people saw themselves as wealthier. We conclude by discussing the role that different evaluations of wealth might play in patterns of conspicuous consumption.

Batista, Mao, and Sussman (2025, Invited 2nd round revision at *JMR*)

This paper examines why consumers simultaneously maintain low-yield liquid assets and high-interest revolving debt, a puzzling financial behavior known as "co-holding." Approximately one in five credit card users in our data exhibit this behavior, effectively paying to borrow back their own deposited money at considerable cost. We propose that co-holding reflects specific spending preferences and money management strategies rather than mere inattention or financial illiteracy. Through analysis of over 3 million transactions from a large retail bank spanning 38 months, we demonstrate that co-holders (n=979) exhibit distinct spending patterns compared to non-co-holders (n=1,328). Specifically, co-holders prefer to pay for everyday expenses using debit cards while reserving credit cards for larger purchases, a pattern robust to controls for demographics, total spending, credit access, and utilization. An online experiment provides causal evidence that inducing preferences for debit spending increases propensity to co-hold. Furthermore, survey data reveals that co-holders hold different beliefs about appropriate uses of credit, perceiving everyday purchases as less appropriate for credit financing. These findings contribute to financial decision-making literature by linking static balance behaviors to dynamic spending patterns and offering insights for financial institutions seeking to design more effective products and interventions. Rather than mere financial mismanagement, co-holding appears to reflect deliberate money management strategies that prioritize payment method preferences despite their substantial costs.

Batista, Mao, Sussman, Mahoney, and Min (2025, Invited 2nd round revision at *PNAS Nexus*)
Why do consumers simultaneously maintain low-yield liquid assets and high-interest revolving debt? This behavior, known as “co-holding,” affects 23% of credit card users in our sample from a major international bank and costs the typical co-holder hundreds of dollars annually in unnecessary interest charges. Our analysis of 38 months of detailed transaction data reveals that co-holding is highly persistent (typical co-holders maintain this behavior for most months observed) and is reflected in spending patterns as well as balances, suggesting deliberate rather than neglectful behavior. We test whether co-holding could be addressed through information disclosure; specifically, whether co-holding is due to inattention or a misunderstanding of the costs. We conducted a large-scale field experiment (n=125,328) providing clear information about co-holding behavior and its costs through the bank's mobile app at moments when customers could easily act. Despite sufficient power to detect small effects, this intervention did not meaningfully affect debt repayment. While customers who received notifications made more frequent payments, they did not reduce their total debt more than the control group. These results challenge the sufficiency of explanations based on limited attention or information gaps and suggest that simple information disclosure, even when carefully designed and delivered through trusted channels, may not effectively address costly financial behaviors.

Batista, Schroeder, Mittal, and Mullainathan (2024)
People do not always say what they mean. In everyday conversations, people regularly sense that they have not fully communicated what they had in mind—a subjective experience we refer to as “misarticulation.” Understanding the nature and causes of misarticulation is crucial, as it affects people’s decisions to communicate. We propose that it is not enough to know something to feel it can be communicated well; rather, misarticulation is also influenced by the way concepts are represented in one’s mind (e.g., the complexity of a concept or the typicality of its constituent objects). Studies 1a-b (n=600) examine when, where, with whom, and why people experience misarticulation. Studies 2-3 (n=1,341) then teach participants concepts of varying complexity and typicality, respectively, before tasking them with articulating the concepts. More complex and atypical concepts elicited stronger misarticulation. Overall, misarticulation is a common, consequential experience influenced by the structure of thought.

Batista and Griffiths (2026)
People increasingly use large language models (LLMs) to explore ideas, gather information, and make sense of the world. In these interactions, they encounter agents that are overly agreeable. We argue that this sycophancy poses a unique epistemic risk to how individuals come to see the world: unlike hallucinations that introduce falsehoods, sycophancy distorts reality by returning responses that are biased to reinforce existing beliefs. We provide a rational analysis of this phenomenon, showing that when a Bayesian agent is provided with data that are sampled based on a current hypothesis the agent becomes increasingly confident about that hypothesis but does not make any progress towards the truth. We test this prediction using a modified Wason 2-4-6 rule discovery task where participants (N=557) interacted with AI agents providing different types of feedback. Unmodified LLM behavior suppressed discovery and inflated confidence comparably to explicitly sycophantic prompting. By contrast, unbiased sampling from the true distribution yielded discovery rates five times higher. These results reveal how sycophantic AI distorts belief, manufacturing certainty where there should be doubt.