



Yale University

Thurman Arnold Project

Digital Platform Regulation Conference
Paper 5

Manipulating, Fast and Slow

May 2022

Jerry Ma

Executive summary

Leading digital platforms rely on a variety of algorithmic optimization techniques to drive activity and growth. These techniques operate both to acquire new users and to induce current users to perform actions desirable to platform owners. Using algorithms, a platform can influence users to spend more time on the platform, engage more actively with the platform’s content and features, contribute more content to the platform (often for free), and produce more revenue for the platform (*e.g.*, by viewing ads).

This article highlights an emerging trend in platforms’ use of algorithms: a shift from “myopic” optimization to “long-horizon” optimization. Myopic optimization maximizes a per-interaction metric—such as the likelihood of a “reshare” or the expected revenue from a given piece of content. In contrast, long-horizon optimization maximizes—for each individual platform-user interaction—the long-term extractable value of the user to the platform. A sufficiently sophisticated long-horizon algorithm practices a form of covert psychology by exploiting any statistically discernible instrument for inducing long-term behavioral modification—including fostering platform addiction—if such a result benefits the platform.

Technology companies’ increasingly ambitious R&D investments in long-horizon optimization must not escape the notice of policymakers. By investing in publicly funded research as to the prevalence and effectiveness of long-horizon optimization algorithms, granting regulators visibility into precisely how each platform trains these algorithms, and imposing a system of context-appropriate guardrails on the use of such algorithms, policymakers can swiftly respond to the critical consequences of long-horizon optimization on user welfare and the broader platform marketplace.

Introduction

The age of digital platforms is coextensive with the age of algorithmic optimization. Platforms such as Twitter, Facebook, YouTube, TikTok, and Instagram employ algorithms to optimize business outcomes ranging from userbase growth to content monetization (e.g., through digital advertising).¹ Large platform companies commonly orient their own R&D divisions toward devising and refining platform-specific algorithms, while startups can make use of Software-as-a-Service (SaaS) offerings that provide slightly less bespoke—but still powerful—algorithmic optimization capabilities.²

Customarily, the workhorse of such techniques are artificial intelligence (AI) and machine learning (ML) models, which learn from data to help platforms automatically make user-facing decisions according to a given objective. Specifically, algorithmic systems that incorporate AI/ML models typically ingest historical data and/or real-time observations and produce quantitative representations thereof. These representations can be used descriptively to characterize known data points, or predictively to forecast unknown outcomes.³ The thousands of employees and billions of capital expenditure dollars directed to AI/ML capabilities are merely a rough indicator of digital platforms' inextricable reliance on AI/ML for algorithmic optimization.⁴

Predictive AI/ML models lend themselves particularly well to algorithmic optimization. As a paradigmatic algorithmic optimization workflow, a platform can specify an “action space” (i.e., a set of

¹ Sharma, Aneesh, Jerry Jiang, Praveen Bommanavar, Brian Larson, and Jimmy Lin. “GraphJet: Real-Time Content Recommendations at Twitter.” *Proceedings of the VLDB Endowment* 9, no. 13 (September 2016): 1281–92; He, Xinran, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, et al. “Practical Lessons from Predicting Clicks on Ads at Facebook.” In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*, 1–9. ADKDD’14. New York, NY, USA: Association for Computing Machinery, 2014; Sarukkai, Ramesh R. “Real-Time User Modeling and Prediction: Examples from YouTube.” In *Proceedings of the 22nd International Conference on World Wide Web*, 775–76. Rio de Janeiro, Brazil: ACM, 2013.

² Amazon Web Services, Inc. “Machine Learning – Amazon Web Services.” Accessed May 23, 2022. <https://aws.amazon.com/sagemaker/>; Google Cloud. “Vertex AI.” Accessed May 23, 2022. <https://cloud.google.com/vertex-ai>.

³ Ma, Alfred C. “Making Data Reports Useful: From Descriptive to Predictive.” *Cureus*, October 12, 2020.

⁴ Moss, Sebastian. “Facebook Plans Huge \$29-34 Billion Capex Spending Spree in 2022, Will Invest in AI, Servers, and Data Centers.” *Data Center Dynamics*, November 1, 2021. <https://www.datacenterdynamics.com/en/news/facebook-plans-huge-29-34-billion-capex-spending-spree-in-2022-will-invest-in-ai-servers-and-data-centers/>; Krauth, Olivia. “The 10 Tech Companies That Have Invested the Most Money in AI.” *TechRepublic*, January 12, 2018. <https://www.techrepublic.com/article/the-10-tech-companies-that-have-invested-the-most-money-in-ai/>; Shead, Sam. “DeepMind A.I. Unit Lost \$649 Million Last Year and Had a \$1.5 Billion Debt Waived by Alphabet.” *CNBC*, December 17, 2020. <https://www.cnbc.com/2020/12/17/deepmind-lost-649-million-and-alphabet-waived-a-1point5-billion-debt.html>.

possible actions the platform can take in each platform-user interaction), build an AI/ML model to predict the likelihood of desirable outcomes across the action space, and then choose the action that can be expected to produce the most desirable outcome. Platforms can mix-and-match predictive models in their overall algorithmic portfolio—for instance, using one model to estimate “likes”, another to estimate “follows”, and so on—toward satisfying a diverse set of business objectives.

As a purely structural matter, the predictive algorithms historically used by digital platforms consider only the immediate platform-user interaction, ignoring most knock-on effects on future states of the world. These classical methods can thus be characterized as “myopic” optimization. However, emerging trends in AI/ML research portend the rise of a new predictive paradigm. This paradigm, which we label “long-horizon” optimization, enables algorithms to encode the impacts that present platform actions will have on future states and outcomes, toward maximizing the overall long-term value of each user to the platform.

The consequences of long-horizon optimization have yet to be formally explored. Yet we contend that they should be deeply worrisome for consumer welfare and platform competition advocates due to the potency of the long-horizon paradigm to modify user behavior. Because long-horizon algorithms can be directed to maximize platform value over an unbounded time period, such algorithms can detect and exploit any behavioral change mechanism that is discernible given the data and the model specification, so long as the resulting behavioral changes benefit the platform. Armed with these algorithms, platforms can extract more content, engagement, and dollars without regard for negative externalities on user well-being, and platforms can make themselves “stickier” to their userbase to mitigate competitive pressures.

We proceed in this article by formalizing the “myopic” versus “long-horizon” temporal taxonomy of algorithmic optimization, then explaining recent developments in AI/ML research that have paved the way for platforms to begin deploying long-horizon optimization algorithms. We then argue that long-horizon optimization, once put into practice within digital platforms, becomes a covert form of statistical psychology and we discuss the implications thereof. Finally, we offer a menu of complementary policy responses to address the serious risks of long-horizon optimization.

Algorithmic optimization: a temporal taxonomy

We begin by precisely defining myopic and long-horizon algorithmic optimization.

A **myopic optimization algorithm** selects, for a single platform-user interaction, an action that maximizes a platform-specified objective function quantifying the value derived by the platform from the interaction itself.⁵

A **long-horizon optimization algorithm** selects, for a single platform-user interaction, an action that maximizes the sum of a platform-specified objective function quantifying the value derived by the platform from the current interaction and *each* future interaction.⁶

A simple example illustrates the distinction in plain terms. Suppose that a digital platform is in the business of showing videos, each of which is monetizable, to its users. The “action” to be taken by the platform is thus the selection of an appropriate video upon user request. A myopic algorithm might choose the video that generates the highest expected revenue-on-view. A long-horizon algorithm, on the other hand, might choose the video that, taking into account future knock-on effects, generates the highest expected revenue for all videos watched by the user over some longer time period.

Then, suppose such an algorithm is choosing a video to show a user between two options: video A, which is typically monetizable for 4 cents, and video B, which is typically monetizable for 2 cents. But suppose that watching video B causes the user to also desire to watch video A with 75% probability (but not vice-versa). A myopic algorithm might select video B to reap 4 cents, while a long-horizon algorithm might select video A to reap 2 cents at first, but with an expected 3 cents of revenue from the next video.

Now observe that the foregoing example depicted a simple, first-order intertemporal dependency (*i.e.*, between one platform-user interaction and the next). Consider that in reality, many higher-order intertemporal dependencies may exist in a video consumption setting—for instance, a music video creating a persistent affinity for other songs from the artist—and the comparative advantage of long-horizon

⁵ Because the outcome of any given interaction is random, maximizing a value typically refers to maximizing the statistical expectation of the value.

⁶ In reality, practical matters often preclude the optimization of the objective function into the indefinite future—for example, a platform might only have a few years of data to work with, or the optimization methods might be too numerically unstable. In such cases, one might impose a reasonable time limit on future interactions to be considered, or use a mathematical technique known as “exponential discounting” for a similar outcome.

optimization from the platform’s perspective is yet more pronounced. And of course, in the general case, platform companies most certainly wish to optimize future business outcomes in their present-day decisions.

From supervised learning to reinforcement learning

Facially, long-horizon algorithmic optimization seems clearly more desirable for platform companies than myopic optimization. Why, then, haven’t long-horizon techniques seen broader adoption within platforms? The answer lies not in business sense but in the raw technology.

For decades, progress in AI/ML has been concentrated within the subfield of “supervised learning”. In supervised learning, an algorithm is trained simply to infer, from a given dataset, a mathematical mapping between each data point and a desired attribute about the data point.⁷ Even with the advent of modern deep learning methods utilizing large neural networks, “supervised learning” (along with related techniques in “semi-supervised” or “self-supervised” learning) has predominated in the list of blockbuster AI headlines that capture public attention.⁸

Supervised learning is most readily usable as a myopic optimization technique. For example, it is straightforward to train a predictive AI/ML model using supervised learning on past data to estimate the probability that a user will click on a given advertisement or give a “like” to a particular post.⁹ While it is theoretically possible to use supervised learning techniques to estimate longer-term outcomes arising from a given interaction (such as lifetime advertisement revenue), the lack of explicit intertemporal structure in such techniques typically renders this task a difficult if not impossible affair.

By contrast, the AI/ML subfield of “reinforcement learning” explicitly accounts for intertemporal dynamics. In reinforcement learning, an algorithm is trained to devise a “policy” (*i.e.*, rule for choosing

⁷ Russell, Stuart J., Peter Norvig, and Ernest Davis. *Artificial Intelligence: A Modern Approach*. 3rd ed. Prentice Hall Series in Artificial Intelligence. Upper Saddle River: Prentice Hall, 2010.

⁸ Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. “ImageNet Classification with Deep Convolutional Neural Networks.” *Communications of the ACM* 60, no. 6 (May 24, 2017): 84–90; Jumper, John, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, et al. “Highly Accurate Protein Structure Prediction with AlphaFold.” *Nature* 596, no. 7873 (August 26, 2021): 583–89; Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. “Attention is All You Need.” *Advances in Neural Information Processing Systems* 30 (2017).

⁹ Cheng, Heng-Tze, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishi Aradhye, Glen Anderson, et al. “Wide & Deep Learning for Recommender Systems.” In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*, 7–10. Boston, MA, USA: ACM, 2016; He et al., “Practical Lessons from Predicting Clicks on Ads at Facebook.”

actions) that maximizes a cumulative future reward.¹⁰ Thus, reinforcement learning is directly aligned with the aims of long-horizon algorithmic optimization.

Because reinforcement learning has historically lagged far behind supervised learning, such techniques were not sufficiently mature for deployment in the platforms that grew to become today's market leaders. But over the past half-decade, reinforcement learning has enjoyed its own robust cadence of breakthrough results. Google DeepMind has arguably formed the vanguard in this R&D revolution, delivering algorithms that produced superhuman performance in the games of Go and StarCraft II.¹¹ OpenAI followed with notable success in the complex team-based game of Dota 2,¹² and Facebook AI Research thereafter in the partial-information environment of no-limit Texas hold 'em.¹³

Although recent progress in reinforcement learning largely originates from artificial settings—such as games or physics simulation engines—these breakthroughs certainly haven't escaped platform companies' notice. For instance, Facebook made a sizable investment starting in 2018 toward translating such breakthroughs into applied settings. The resulting reinforcement learning software capabilities have subsequently been deployed in product-user interactions such as mobile push notifications.¹⁴ Similar investments by companies such as YouTube and TikTok have also recently materialized in real-world products toward engagement and revenue maximization.¹⁵

¹⁰ Reinforcement learning is closely tied to optimal control theory as used in mathematics, economics, engineering, and other technical fields. Common formulations of reinforcement learning problems define the task as solving the Bellman equation or Hamilton-Jacobi-Bellman equation for a given Markov decision process.

¹¹ Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, et al. "Mastering the Game of Go with Deep Neural Networks and Tree Search." *Nature* 529, no. 7587 (January 28, 2016): 484–89; Vinyals, Oriol, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, et al. "Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement Learning." *Nature* 575, no. 7782 (November 14, 2019): 350–54.

¹² Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, et al. "Dota 2 with Large Scale Deep Reinforcement Learning." *ArXiv*:1912.06680, December 13, 2019.

¹³ Brown, Noam, and Tuomas Sandholm. "Superhuman AI for Multiplayer Poker." *Science* 365, no. 6456 (August 30, 2019): 885–90; Brown, Noam, Anton Bakhtin, Adam Lerer, and Qucheng Gong. "Combining Deep Reinforcement Learning and Search for Imperfect-Information Games." In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 17057–69. NeurIPS'20. Red Hook, NY, USA: Curran Associates Inc., 2020.

¹⁴ Gauci, Jason, Edoardo Conti, Yitao Liang, Kittipat Virochsiri, Yuchen He, Zachary Kaden, Vivek Narayanan, Xiaohui Ye, Zhengxing Chen, and Scott Fujimoto. "Horizon: Facebook's Open Source Applied Reinforcement Learning Platform." *ArXiv*:1811.00260, September 4, 2019; Liu, Yang, Zhengxing Chen, Kittipat Virochsiri, Juan Wang, Jiahao Wu, and Feng Liang. 2021. "Reinforcement Learning-Based Product Delivery Frequency Control". *Proceedings of the AAAI Conference on Artificial Intelligence* 35 (17):15355-61.

¹⁵ Chen, Minmin, Alex Beutel, Paul Covington, Sagar Jain, Francois Belletti, and Ed H. Chi. "Top-K Off-Policy Correction for a REINFORCE Recommender System." In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 456–64. Melbourne, VIC, Australia: ACM, 2019; Zhao, Xiangyu,

Future progress in reinforcement learning research will only bring about more opportunities for platforms to deploy long-horizon optimization algorithms. And so long as business pressures continue to incentivize platform companies to optimize future engagement and revenue metrics, platforms will rely more heavily on long-horizon optimization algorithms as the underlying reinforcement learning technology matures.

Long-horizon optimization as covert statistical psychology

Digital platform companies have long recognized the importance of understanding user psychology in operating their platforms. Leaders in the sector routinely hire teams of behavioral scientists to develop features that influence users for private objectives.¹⁶ Although “dark nudges” or “dark patterns”—product features that exploit idiosyncrasies in consumer behavior—are frequently deployed across diverse commercial sectors,¹⁷ the pervasiveness of social media and digital content in daily life renders behavioral product design both particularly lucrative and worrisome in the digital platform context.¹⁸

Yet behaviorally-informed product design still retains a modicum of interpretability. That is, a product intentionally designed to exploit consumer behavior relies on particular idiosyncrasies derived from behavioral science research. Once such design choices are uncovered, regulators and the public can make their own judgments as to the propriety of the behavioral manipulations employed by each platform, and platforms can be held to account accordingly.

This paradigm is flipped on its head in the world of algorithmic optimization. Because algorithmic optimization techniques rely on mathematics rather than on behavioral science, it is facially unapparent that such techniques can do anything to exploit idiosyncrasies in user behavior. But reality undermines this

Xudong Zheng, Xiwang Yang, Xiaobing Liu, and Jiliang Tang. “Jointly Learning to Recommend and Advertise.” In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 3319–27. Virtual Event: ACM, 2020.

¹⁶ Lieber, Chavie. “Psychologists Are Speaking out against Tech Companies That Use Psychology to Lure Kids In.” *Vox*, August 8, 2018. <https://www.vox.com/2018/8/8/17664580/persuasive-technology-psychology>.

¹⁷ Newall, Philip W. S. “Dark Nudges in Gambling.” *Addiction Research & Theory* 27, no. 2 (March 4, 2019): 65–67; Petticrew, Mark, Nason Maani, Luisa Pettigrew, Harry Rutter, and May Ci Van Schalkwyk. “Dark Nudges and Sludge in Big Alcohol: Behavioral Economics, Cognitive Biases, and Alcohol Industry Corporate Social Responsibility.” *The Milbank Quarterly* 98, no. 4 (December 2020): 1290–1328.

¹⁸ Stigler Committee on Digital Platforms. “Final Report,” September 2019. <https://www.chicagobooth.edu/research/stigler/news-and-media/committee-on-digital-platforms-final-report>; United States House of Representatives, Committee on the Judiciary, Subcommittee on Antitrust, Commercial and Administrative Law. “Investigation of Competition in Digital Markets: Majority Staff Report and Recommendations.” <https://purl.fdlp.gov/GPO/gpo145949>.

seeming behavioral neutrality. A good optimization algorithm, endowed with a given quantitative objective, will tease out any discernible mechanism for maximizing that objective from the data. And if the most powerful mechanisms happen to manipulate user behavior, then algorithms will naturally morph into automated behavioral manipulation engines.

Three major implications follow from this observation. First, **platforms' optimization algorithms can be characterized as covert statistical psychology**. These algorithms are *statistical* psychology because they venture beyond the frontiers of behavioral science research. Without any domain knowledge or psychological expertise, these algorithms can simply mine large datasets for promising behavioral quirks that probabilistically benefit platforms' bottom lines.

Moreover, they are *covert* because the underlying AI/ML models largely operate in the domain of opaque mathematical vector spaces. Thus, unlike human behavioral scientists, these algorithms simply cannot provide an account of *how* they are manipulating user behavior. And because users are already inattentive to habit formation in the context of digital platforms,¹⁹ the covert nature of algorithmic behavioral modification buries it even further within the recesses of public consciousness.

Second, because these algorithms perform covert statistical psychology rather than overt human-implemented psychology, **platform companies that employ algorithmic optimization enjoy an effective shield from public scrutiny**. In deploying optimization algorithms, platform companies can more easily deny behavioral effects of platform features, profess *ex ante* ignorance after such denials are no longer plausible, and plead lack of intent when governments and civil society attempt to hold them to account.

Consider the example offered by late 20th-century tobacco companies, which knew that their products addicted consumers through nicotine and nonetheless deliberately designed tobacco products to maximize nicotine delivery.²⁰ The ensuing regulatory crackdown on tobacco products was arguably enabled by factual findings of the products' addictiveness, the specific dependence-inducing chemical pathway (*i.e.*,

¹⁹ Allcott, Hunt, Matthew Gentzkow, and Lena Song. "Digital Addiction." Cambridge, MA: National Bureau of Economic Research, June 2021.

²⁰ Kessler, D. A. "The Control and Manipulation of Nicotine in Cigarettes." *Tobacco Control* 3, no. 4 (December 1, 1994): 362–69; Hurt, Richard D., and Channing R. Robertson. "Prying Open the Door to the Tobacco Industry's Secrets About Nicotine: The Minnesota Tobacco Trial." *Journal of the American Medical Association* 280, no. 13 (October 7, 1998): 1173.

nicotine), and the industry's intent to exploit this pathway toward long-term, profit-driven behavioral change.²¹

Now imagine that regulators and the public lacked any causally convincing evidence of tobacco's addictiveness, any idea of the specific mechanism underlying said addictiveness, or any indicia of an intent to chemically addict consumers. It is unlikely that the force of public opinion and regulatory response would have approached anything resembling that of the 1990's Big Tobacco crackdown.

This is precisely the scenario in which we find ourselves today thanks to algorithmic optimization. Recent literature has offered descriptive evidence of digital platforms' unanticipated behavioral and addictive effects.²² But the frequent absence of discretely identifiable behavioral manipulations—or strong evidence of intent to manipulate—operates both to diffuse public outrage and to stymie regulators who often require more than descriptive evidence alone to act. As a result, platform companies can employ a convenient narrative—“we're just doing math here”—as a defense against unaccountable platform outcomes.

Finally, **the emergence of long-horizon optimization greatly amplifies the effectiveness—and the consequent concerns—of algorithms as behavioral modification tools.** In the world of myopic optimization, algorithmic decision-making is cabined within the context of a specific platform-user interaction. And while even myopic algorithms are prone to exploit idiosyncrasies in user behavior, the fact that such algorithms optimize over the short term partially mitigates the risk and severity of behavioral modification.

With long-horizon algorithms, persistent behavioral modification is far likelier to arise as an explicit mechanism of algorithms rather than as a side effect. For instance, should an algorithm infer that addicting users to the platform maximizes its optimization objective, the algorithm will exploit any statistically discernible mechanism for inducing that addiction. Simply by reducing each user to an equation

²¹ Kessler, David A. *A Question of Intent: A Great American Battle with a Deadly Industry*. New York: Public Affairs, 2001.

²² Levounis, Petros, James Sherer, and American Psychiatric Association Publishing, eds. *Technological Addictions*. First edition. Washington, DC: American Psychiatric Association Publishing, 2022; Montag, Christian, Bernd Lachmann, Marc Herrlich, and Katharina Zweig. “Addictive Features of Social Media/Messenger Platforms and Freemium Games against the Background of Psychological and Economic Theories.” *International Journal of Environmental Research and Public Health* 16, no. 14 (July 23, 2019): 2612.

to be maximized, a long-horizon algorithm will figure out the messy details of how to keep eyeballs glued, fingers tapping, and whatever other actions might benefit the platform.

Given the already-serious unintended consequences of myopic optimization—*e.g.*, hate speech, polarization, and youth self-esteem issues—²³ there is ample reason to fear even graver results when addiction and other persistent behavioral changes are instrumentally pursued by long-horizon optimization algorithms. And because any behavioral changes would be covert and indiscrete, companies that turn to long-horizon algorithms can continue their attempts to minimize the resulting scrutiny.

Policy responses to long-horizon optimization

As technology companies invest ever-greater sums in the research and computational infrastructure underlying long-horizon optimization, policymakers and regulators should prepare to vigorously defend consumers. A sound policy response should prioritize three prongs: research and analysis, transparency, and guardrails.

At the outset, regulators must understand the current maturity, prevalence, and effectiveness of long-horizon optimization—a tall order given the dearth of publicly available information in this area. Therefore, policymakers must invest in publicly funded research and analysis that will help answer crucial questions. For example, how widespread are long-horizon optimization algorithms within digital platforms? How effectively do modern reinforcement learning algorithms perform long-horizon optimization in the digital platform context, and how quickly are these algorithms improving over time? How can regulators uncover the behavioral modification channels exploited by long-horizon algorithms, when said channels may be intricate statistical dependencies that defy verbal explanation? Because such efforts will require a mix of technical expertise and policy analysis acumen, policymakers should facilitate these research activities through collaborations between regulatory agencies and academia.

²³ Frenkel, Sheera, and Davey Alba. “In India, Facebook Grapples With an Amplified Version of Its Problems.” *The New York Times*, October 23, 2021. <https://www.nytimes.com/2021/10/23/technology/facebook-india-misinformation.html>; Barrett, Paul M., Justin Hendrix, and J. Grant Sims. “Fueling the Fire: How Social Media Intensifies U.S. Political Polarization—And What Can Be Done About It.” NYU Stern Center for Business and Human Rights. September 2021. <https://bhr.stern.nyu.edu/polarization-report-page>. Tatum Hunter. “For Teens, Navigating the Mental Health Pitfalls of Instagram Is Part of Everyday Life.” *The Washington Post*, October 21, 2021. <https://www.washingtonpost.com/technology/2021/10/21/teens-instagram-feed-mental-health/>.

After developing a strong general understanding of long-horizon optimization, regulators must then turn to specific instances in which these techniques are being deployed. But regulators currently lack the ability to examine these instances because the internal technology simply isn't accessible to agencies. Prior work has proposed various transparency standards for algorithms, datasets, and other components routinely used in algorithmic optimization.²⁴ But because the optimization objective itself determines, in large part, the types of behavioral modification a long-horizon algorithm might induce in users, policymakers should consider requiring platform companies to make the mathematical objective functions used in long-horizon optimization available to regulators. Such objective functions should be accompanied by any data science artifacts—for example, code notebooks, quantitative analyses, and so on—generated by the company in determining whether to employ that objective function.

Finally and most crucially, policymakers must construct proper regulatory guardrails to control any negative impacts of long-horizon optimization. The Federal Trade Commission (FTC) likely possesses authority in this area already, as a long-horizon optimization algorithm that covertly induces persistent behavioral change could almost certainly be construed as an “unfair or deceptive act or practice” under Section 5 of the FTC Act.²⁵ Moreover, such an algorithm could plausibly become an “unfair method of competition” under Section 5 as well—algorithm-facilitated platform addiction operates to the detriment of competitors and undercuts the consumer choice-oriented arguments that large platforms have historically marshaled as defenses to antitrust enforcement.²⁶

Nonetheless, the FTC may not be the optimal regulator in this area. Beyond the severe resource constraints and already-broad portfolio of the agency,²⁷ its usual toolkit of investigatory and litigation

²⁴ Mitchell, Margaret, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. “Model Cards for Model Reporting.” In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–29. Atlanta GA USA: ACM, 2019; Pushkarna, Mahima, Andrew Zaldivar, and Oddur Kjartansson. “Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI.” *ArXiv:2204.01075*, April 3, 2022; United Kingdom Cabinet Office, Central Digital and Data Office. “Algorithmic Transparency Standard.” November 29, 2021. <https://www.gov.uk/government/collections/algorithmic-transparency-standard>.

²⁵ 15 U.S.C. § 45.

²⁶ Google LLC. “Answer to Amended Complaint.” *United States of America v. Google LLC (1:20-cv-03010)*. January 29, 2021; Facebook, Inc. “Memorandum in Support of Facebook, Inc.'s Motion to Dismiss FTC's Complaint.” *Federal Trade Commission v. Meta Platforms, Inc. (1:20-cv-03590)*, March 10, 2021.

²⁷ Christine S. Wilson. “Governing Is Hard: Antitrust Enforcement in the First Year of the Biden Administration.” Presented at the Mercatus Antitrust Forum, January 26, 2022. https://www.ftc.gov/system/files/documents/public_statements/1600479/governing_is_hard_antitrust_enforcement_i

authority likely falls short of what is needed. FTC investigations can take years even before any potential suit, while platforms' optimization algorithms are released, refined, and retired on a much more frequent cadence.²⁸ Moreover, a regulatory regime that relies on proactive investigation alone cannot hope to comprehensively protect consumers, who frequent an increasingly diverse bevy of digital platforms.

Therefore, policymakers might look toward instituting a new algorithmic regulatory agency. Such an agency could be endowed with the traditional investigatory and litigation powers that characterize a consumer protection regulator. But such an agency might also assume a preapproval role—*i.e.*, a requirement for platform companies to report proposed algorithm deployments to the agency and await regulatory approval thereof. This regime could be modeled after the Hart-Scott-Rodino (HSR) Act, which prescribes an analogous requirement for mergers and acquisitions. In place of transaction size thresholds, policymakers could define criteria for mandatory preapproval—such as the size of the platform, the number of users affected by the algorithm, past instances of persistent behavioral modification involving the platform's algorithms, or the algorithm's use of a long-horizon optimization objective. And similar to the procedures in the HSR Act, the regulatory agency would be able to allow the algorithm's deployment, to negotiate an agreement with the platform company that mitigates consumer welfare concerns, or to initiate administrative proceedings to block deployment on consumer welfare grounds.

Policymakers could look to other mechanisms as either alternatives or complements to this preapproval regime—for instance, blanket prohibitions on certain types of long-horizon optimization algorithms as applied to consumers, statutory causes of action for algorithm-induced behavioral modification, or stringent data portability standards to enable the entry of competitors less reliant on algorithmic optimization. But however policymakers respond to long-horizon optimization, they must

[n the first year of the biden administration 0.pdf](#); Federal Trade Commission. “FTC Testifies Before House Energy and Commerce Subcommittee on Legislation to Modify the Commission’s Authority and Address Challenges Facing the Agency,” July 28, 2021. <http://www.ftc.gov/news-events/news/press-releases/2021/07/ftc-testifies-house-energy-commerce-subcommittee-legislation-modify-commissions-authority-address>.

²⁸ David Willman. “PGA Outclubs FTC in Antitrust Fight.” *Los Angeles Times*, October 22, 1995. <https://www.latimes.com/archives/la-xpm-1995-10-22-fi-59876-story.html>; Federal Trade Commission. “FTC Enforcement Action Leads U.S. Dept. of Education to Forgive \$71.7 Million in Loans for Students Deceived by DeVry University,” February 16, 2022. <http://www.ftc.gov/news-events/news/press-releases/2022/02/ftc-enforcement-action-leads-us-dept-education-forgive-717-million-loans-students-deceived-devry>.

ensure that the guardrails governing the use of these powerful and opaque techniques are strong enough to comprehensively protect digital consumers.

Conclusion

Algorithmic accountability has recently emerged as a focal point for governments and civil society. But policy discussions have largely overlooked the temporal spectrum of algorithmic optimization. All the while, tech industry juggernauts are deploying researchers and dollars toward extending this spectrum through developing long-horizon optimization techniques such as reinforcement learning. The resulting consequences to consumer welfare and platform competition—including via persistent behavioral modification—must not escape scrutiny.

To be clear, there are no indications of overt malintent from large platform companies, who may deploy such algorithms in a method-agnostic effort to bolster the bottom line. The scientists and engineers responsible for developing long-horizon algorithms are not themselves encoding behavioral manipulation techniques into the software they write. But once given a long-horizon optimization objective, these algorithms will nonetheless seize on whatever dependencies maximize said objective—including those that give rise to addiction, habit formation, and other persistent behavioral impacts to users.

Thus, governments must confront the incoming tide of algorithm-induced behavioral modification with the same gravity they mustered toward the tobacco epidemic and other major consumer protection crises. Only through descriptive research, regulatory visibility, and strong guardrails can policymakers ensure that these technological innovations are deployed in ways that promote rather than impair consumer welfare.