Case Study: The Distilling of a Biased Algorithmic Decision System through a Business Lens

Merve Hickok (Alethicist.org)

Colleen Dorsey (University of St Thomas School of Law)

Tim O'Brien (Independent Researcher)

Dorothea Baur (Baur Consulting)

Katrina Ingram (Ethically Aligned AI)

Chhavi Chauhan (The American Society for Investigative Pathology)

Attlee Gamundani (UN University Institute in Macau)

2022

Contents

Introduction	3
Introduction to the case: Impact Pro > Optum > UnitedHealth Group	7
Defining the problem	9
Risks to the business	13
Diversity of teams & stakeholders	15
Assurance & governance	18
Data, processes & model details	22
Deployment	24
Monitoring, feedback, calibration / improvement	27
Know the context of your data historical / structural bias in healthcare	29
Conclusion	32
Information on Optum	34
Acknowledgements	34

All rights reserved by the authors.

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

The Distilling of a Biased Algorithmic Decision System Through a Business Lens

Who should read this case study?

C-suite and Board level executives - to understand bias, harms, and consequences in algorithmic decision systems; to enable responsible and sustainable development and implementation within their organizations; to create necessary governance mechanisms for oversight

Policymakers - to understand existing biases in algorithmic decision systems; to establish safeguards and oversight that shapes better practices, and to implement policy to mitigate these biases

Development & Implementation teams - to implement the guidance throughout the lifecycle of these systems; to design and deploy scalable, sustainable, and better products that serve everyone

Introduction

Technological advances embedded within algorithmic decision systems¹ are being deployed every minute of every day. Some of these investments are unquestionably worthwhile, while some others prioritize commercialization of technology ahead of societal impact. PwC estimates that artificial intelligence (AI) could boost global GDP by 14% by 2030, or up to US\$ 15.7 trillion.² Tempering these great economic booms, however, is the impact these systems have on our lives in ways which have profound ethical consequences most of the general population could never have imagined. A few examples:

- Black men falsely arrested based on facial recognition systems^{3 4}
- Driverless cars striking and killing pedestrians⁵
- An algorithm blocking kidney transplants to Black patients⁶
- Fake “bots” on websites inducing users to make fraudulent purchases.⁷
- Algorithmic recruiting tools in use biased against female applicants⁸
- State agencies utilizing matching algorithms that automatically and inaccurately disqualified individuals for food assistance⁹

¹ Note: Choice of the word ‘algorithmic decision systems’ by the authors is intentional. It covers not only the tool itself, but also the data used, and the humans involved in the data and design practices, as well as the use of these systems. The word choice also means that although ‘algorithmic’ first brings to mind an AI system, one can apply the same questions to the design and use of legacy software systems not powered by AI.

² Rao, Anand S., Verweij, Gerard (2017) Sizing the Price. PWC. <https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html>

³ Allyn, Bobby (2020) ‘The Computer Got It Wrong’: How Facial Recognition Led to False Arrest of Black Man. NPR. <https://www.npr.org/2020/06/24/882683463/the-computer-got-it-wrong-how-facial-recognition-led-to-a-false-arrest-in-michig>

⁴ Tate Ryan-Mosley (2021) The new lawsuit that shows facial recognition is officially a civil rights issue. MIT Technology Review. <https://www.technologyreview.com/2021/04/14/1022676/robert-williams-facial-recognition-lawsuit-aclu-detroit-police/>

⁵ Gonzales, Richard (2019) Feds Say Self-Driving Uber SUV Did Not Recognize Jaywalking Pedestrian In Fatal Crash. NPR. <https://www.npr.org/2019/11/07/777438412/feds-say-self-driving-uber-suv-did-not-recognize-jaywalking-pedestrian-in-fatal>

⁶ Simonite, Tim, (2020), How an Algorithm Blocked Kidney Transplant to Black Patients, Wired. <https://www.wired.com/story/how-algorithm-blocked-kidney-transplants-black-patients/>.

⁷ Farinelli, Bruno (2018) Fraud Bots: Is Your Business at Risk? ClearSale. <https://blog.clear.sale/fraud-bots-is-your-business-at-risk>

⁸ Dastin, Jeffrey (2018) Amazon scraps secret AI recruiting tool that showed bias against women. Reuters <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>

⁹ Gilman, Michele (2020) AI algorithms intended to root out welfare fraud often end up punishing the poor instead. The Conversation. <https://theconversation.com/ai-algorithms-intended-to-root-out-welfare-fraud-often-end-up-punishing-the-poor-instead-131625>

- Credit card approval algorithm biased against women¹⁰
- State falsely charging thousands with unemployment fraud and took millions from them¹¹
- Healthcare algorithm wrongly terminating home care and community integration funds for people with severe disabilities¹²

There are countless more examples of how ubiquitous algorithmic decision systems have become in our world and how they can be inappropriately utilized.¹³ It is now time to take action to address these issues. But how do we get the business community, the different industries across verticals, to consider the ethical ramifications of technology? This question applies equally whether an organization is developing, or utilizing, these systems. Some may not be aware of the high-level principles and academic research which has been conducted and shared concerning these subjects. Some may be only vaguely aware, but unsure what actions to take as a result of these findings. And, certainly, there are some who operate organizations and industries which are in lockstep with today's technological advances and the ethical issues they raise. However, for the lion's share of business leaders, this is uncharted territory.

The authors of this case study, a group of AI ethicists, joined forces to analyze a real-world case from the healthcare sector to demonstrate design and governance shortfalls of an algorithmic system through its lifecycle. The case study provides guidance on responsible and sustainable deployment practices applicable across industries. What follows is an illustration of how the business community can utilize a case study, which goes beyond the news headlines and includes its attendant nuances, to explain what went wrong across the lifecycle of an algorithmic decision-making system, why it matters, and what could have been done to avoid or mitigate the risks and harms.

¹⁰ Justin Herrick (2019) Apple Card: Viral Tweet Leads to Gender Bias Probe, <https://www.pcmag.com/news/apple-card-viral-tweet-leads-to-gender-bias-probe>.

¹¹ Wykstra, Stephanie (2020) Government's Use of Algorithm Serves Up False Fraud Charges. Undark. <https://undark.org/2020/06/01/michigan-unemployment-fraud-algorithm/>

¹² Lecher, Colin (2018) What happens when an algorithm cuts your healthcare. The Verge. <https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy>

¹³ West, S. M., Whittaker, M. & Crawford, K. (2019) Discriminating systems: Gender, race, and power in AI. AI Now Institute. <https://ainowinstitute.org/discriminatingsystems.pdf>

The systematic dissection of this case can be applied to different systems across different domains like employment, credit scoring, housing, education, criminal justice, and many others. *Every business, in every industry, automating and streamlining processes and core functions through digital transformation needs to implement new accountability and governance mechanisms and invest in the inevitable culture change necessary.*¹⁴ Data ethics, bias, discriminatory harms are not for tech industry only, but for every industry and its respective regulator.

Figure 1: Project Lifecycle¹⁵

¹⁴ O'Brien, Tim (2021) Digital Business, Regulated Business. *Tech & Policy*.
<https://timobrien.substack.com/p/digital-business-regulated-business>

¹⁵ The Alan Turing Institute (2021) Human Rights, Democracy, and the Rule of Law Assurance Framework for AI Systems: A proposal prepared for the Council of Europe's Ad hoc Committee on Artificial Intelligence. <https://rm.coe.int/huderaf-coe-final-1-2752-6741-5300-v-1/1680a3f688>

The business case: Why should businesses care about AI ethics?

The severity of the consequences of improperly designed or deployed algorithmic systems is quite clear. Citizens, consumers, users, employees, and even lawmakers have a clear stake in 'AI ethics' but what about companies? What do failures in the ethical development and deployment mean for them, both short term and long term?

Taking AI ethics (understood as the field of applied ethics covering all ethical questions emerging in the context of designing, developing, deploying, using, and monitoring AI), seriously is by no means an entirely altruistic expectation. Most evidently, companies can suffer severe reputational damage¹⁶ when developing or deploying AI with harmful consequences for their stakeholders. As public awareness and pressure increases, companies cannot afford anymore to ignore ethical issues surrounding AI. Neglecting these issues in the design and deployment of AI negatively impacts their bottom line. Boycotts, fines, lawsuits, user, and employee dissatisfaction¹⁷ are costly. As investors are increasingly attuned to non-financial matters, typically summarized by the acronym ESG (environmental, social, governance), we can expect an increase in shareholder activism that also includes questions surrounding AI ethics. Amazon, for example, has experienced this investor pressure regarding their facial recognition system, Rekognition.¹⁸ Ahead of their annual meeting in May 2019, shareholders introduced two proposals on facial recognition for a vote. One asked the company to prohibit sales of Rekognition to government agencies unless its board concluded that the technology did not facilitate human rights violations. The other asked the company to commission an independent report examining the extent to which Rekognition threatened civil, human and privacy rights, and the company's finances.¹⁹ The company ended up putting a moratorium on the use of Rekognition which is still in place today.²⁰ This example illustrates that AI ethics issues are directly 'material' for business – and increasingly important given the present focus on ESG.

Materiality is a key concept in the ESG debate that guides the focus on non-financial matters. In financial reporting, materiality describes “a threshold for influencing the economic decisions of those using an organization's financial statements, investors in

¹⁶ Bonime–Blanc, Andrea (2018) AI and reputational risk: An ESG Perspective. Risk Management. <https://gecrist.com/wp-content/uploads/2018/02/ABonimeBlanc-AI-ESG-and-RepRisk-Ethical-Boardroom-Winter-20182.pdf>

¹⁷ Metz, Rachel (2021) How one employee's exit shook Google and the AI industry. CNN. <https://edition.cnn.com/2021/03/11/tech/google-ai-ethics-future/index.html>

¹⁸ Singer, Natasha (2019) Amazon Faces Investor Pressure Over Facial Recognition. NY Times. <https://www.nytimes.com/2019/05/20/technology/amazon-facial-recognition.html>

¹⁹ *Id.*

²⁰ Dastin, Jeffrey (May 18, 2021), Amazon extends moratorium on police use of facial recognition software, Reuters <https://www.reuters.com/technology/exclusive-amazon-extends-moratorium-police-use-facial-recognition-software-2021-05-18/>

particular.”²¹ In the ESG debate, materiality additionally contains a non-financial dimension, namely the range of a company’s impacts and stakeholders. Materiality builds a bridge between non-financial issues and business strategy. ESG investors expect companies to focus on material issues as part of their corporate responsibility. They want to make sure that companies are aware of the impacts of their activities on stakeholders and their implications for their bottom line. Companies are well-advised to include issues relating to AI ethics in their materiality assessments to be ready to respond to the increasing number of ESG-oriented shareholders.

When done appropriately, embedding AI ethics is not just a means to avoid damages, but to gain advantages.²² An effective and credible AI ethics strategy that takes into account the needs, rights, and expectations of various stakeholders, including investors and customers, can improve the quality of products, enhance their longevity and allow adaptation at scale, strengthen trust in stakeholder relations, increase loyalty among employees, and thus lead to a competitive advantage over ‘ethically blind’ competitors.

Introduction to the case: Impact Pro > Optum > UnitedHealth Group

The healthcare sector sits on a mine of data, making it one of the most lucrative fields for big data–based analytics. However, the remunerative healthcare sector is particularly sensitive to the quality of data and design decisions, as well as data privacy and safety, making it paramount for all stakeholders to safely and ethically develop, deploy and implement the algorithmic tools. Otherwise, these systems can have detrimental effects on the life, wellbeing, and safety of patients. The case study discussed here covers Impact Pro, a data analytics algorithm sold by the UnitedHealth Group’s subsidiary OptumInsight (Optum). Impact Pro is a claims-based risk stratification algorithmic tool that assigns “risk scores” for patients and flags the patients who are at high risk for specialized interventions²³ (Optum).

In 2019, a group of researchers from Berkeley’s School of Public Health, Boston’s Brigham and Women’s Hospital, Massachusetts General Hospital, and the University of Chicago

²¹ Global Reporting Initiative (2018) GRI 101: Foundation 2016
<https://www.globalreporting.org/media/55yhvety/gri-101-foundation-2016.pdf?page=23>

²² The Economist Intelligence Unit (2020) Staying Ahead of the Curve, The business case for responsible AI, The Economist Intelligence <https://www.eiu.com/n/staying-ahead-of-the-curve-the-business-case-for-responsible-ai/>

²³ Impact Pro® for care management (2012) OptumInsight.
https://www.optum.com.br/content/dam/optum/resources/productSheets/Impact_Pro_for_Care_Management_ps_06_2012.pdf

Business School published research²⁴ in the journal *Science*, about Impact Pro.²⁵ Usually, commercial algorithms have proprietary code that restricts the review of the data and model utilized. Obermeyer et al did not have access to the Impact Pro's model but were able to evaluate the dataset²⁶ of a single academic hospital (Partners Healthcare²⁷) as it used Impact Pro. In the health system, "risk scores are generated for each patient during the enrollment period for the system's care management program...Those above the 55th percentile are referred to their primary care physician, who is provided with contextual data about the patients and asked to consider whether they would benefit from program enrollment."²⁸ The research team had access to the inputs, outputs, and function of the model, or in other words, what the model was trying to optimize in that health system. The details give a great and rare opportunity to systematically review this algorithmic decision system as a case from an AI ethics perspective and suggest better practices to developers and decision-makers involved in high-stake situations.

Obermeyer *et al* showed that this widely used algorithm, affecting millions of patients, exhibited significant racial bias: the risk scores produced by the system were putting white patients higher in the queue for specialized interventions than Black patients who were equally sick. Overall, of the patients identified by the algorithm as needing more care, only 18% were Black. If the algorithm were to reflect the true proportion of the sickest patients, Black patients would have made up about 46% of the patients identified instead. In other words, correcting the bias would more than double the number of Black patients flagged as at risk of complicated medical needs within the health system.²⁹

The algorithm was trained on cost data (i.e., dollars spent on healthcare) rather than the health and underlying physiology of the patients. The assumption that cost could be used as a proxy for health need meant that the model valued the people who historically had access to healthcare more than those people who actually needed care. At any high-risk level, Black patients were systemically sicker than white patients.

In the healthcare domain, the use of algorithms has been on an upwards trend. Healthcare data is generated every second around the country. Patient files, laboratory results,

²⁴ Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science (New York, N.Y.)*, 366(6464), 447–453
<https://doi.org/10.1126/science.aax2342> ; and Supplementary Materials
<https://www.science.org/doi/suppl/10.1126/science.aax2342>

²⁵ Bgmanastiri (2019) Pro Case Study: Integrating claims and health risk assessment data to measure risk in a dual eligible population.
<https://bgmanastiri.info/content/dam/optum/resources/caseStudies/PRO%20Case%20Study.pdf>

²⁶ Synthetic Data and Code for Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations (2020) <https://gitlab.com/labsysmed/dissecting-bias/-/blob/master/ReadMe.md>

²⁷ Obermeyer, Ziad. et al (2019)

²⁸ Obermeyer, Ziad. et al (2019)

²⁹ Id.

medications, notes from patient interactions, all the costs associated with treatments, and insurance claims are only a few examples of inflows into healthcare datasets. It is forecasted that almost one-third of the world's stored data is healthcare-related.³⁰ Increasing access to health data in the private sector powers increasingly capable algorithmic systems. Today we see algorithms being tested in the diagnosis of certain diseases, suggestion and analysis of treatments, review of medical literature³¹, predicting protein folding and genetic makeup³², as well as predictions of healthcare costs.³³ In a landscape where healthcare costs are one of the highest expenses for both patients and insurance companies, it is no surprise that both hospitals and insurance companies try to make the most out of the data at hand.

Defining the problem

One of the most crucial and first steps in developing any system is to define the problem to solve and frame the question correctly. How one frames a question impacts what kind of data are to be used, how the outcomes are measured, and what is optimized (objective function) in the system.

A few years ago, the Department of Health & Human Services (HHS) started transitioning from traditional fee-for-service Medicare payments to value-based care in which the central goal was how healthcare providers can improve patient outcomes and quality of care. Success could be measured by criteria like improving preventative care or reducing hospital readmissions. Another goal was to hold healthcare providers accountable by requiring them to report on data-driven decisions, track and demonstrate impact. The “Affordable Care Act created financial incentives for health systems to keep people well instead of waiting to treat them when they got sick”.³⁴ Hospitals and insurance companies now enter into agreements where hospitals are either paid a flat annual fee for each patient (in contrast to the traditional “fee-for-service” model where the more care hospitals provided, the more they were paid), or receive end-of-year monetary adjustments relative to negotiated per-patient cost targets.³⁵ If hospitals can identify patients who are at greater risk for serious illness and provide targeted interventions to keep them healthier, they also contain costs. The incentives of both the patients and hospitals are therefore aligned. In

³⁰ Fry, Erika, Mukherjee, Sy (2018) Tech's Next Big Wave: Big Data Meets Biology. Fortune. <https://fortune.com/2018/03/19/big-data-digital-health-tech/>

³¹ Secinaro, S., Calandra, D., Secinaro, A. *et al.* (2021) The role of artificial intelligence in healthcare: a structured literature review. *BMC Med Inform Decis Mak* 21, 125. <https://doi.org/10.1186/s12911-021-01488-9>

³² Senior, A.W., Evans, R., Jumper, J. *et al.* (2020) Improved protein structure prediction using potentials from deep learning. *Nature* 577, 706–710. <https://doi.org/10.1038/s41586-019-1923-7>

³³ Wolff, J., Pauling, J., Keck, A., & Baumbach, J. (2020) The Economic Impact of Artificial Intelligence in Health Care: Systematic Review. *Journal of medical Internet research*, 22(2), e16866. <https://doi.org/10.2196/16866>

³⁴ Johnson, Carolyn Y (2019)

³⁵ Obermeyer, Ziad. *et al* (2019)

healthcare, the overarching goal for providers, as well as for every other stakeholder, must be improving value for patients, where value is defined as the health outcomes achieved that matter to patients relative to the cost of achieving those outcomes.³⁶

Currently, several predictive algorithms³⁷ impacting approximately 200 million patients across the United States are in use, with different features and assumptions but practically trying to predict the same risk. These include, but are not limited to, ACG System from Johns Hopkins University; Impact Pro, provided by Optum; Clinical Risk Groups (CRG) from 3M; DxCG Intelligence from Verisk Health; HHS-HCC from the CMS; Milliman Advanced Risk Adjusters (MARA) from Milliman; the Prospective Cost of Care Model, from SCIO Health Analytics; the Truven Cost of Care Model, from IBM's Truven Health Analytics; and Wakely Risk Assessment Model from Wakely Consulting Group.³⁸ By 2021, Gartner predicted that 75% of healthcare delivery organizations (HDOs) would have invested in an AI capability that is explicitly improving either operational performance or clinical outcomes.³⁹

Impact Pro, as described in Optum's website, provides customers "analytical foundation to identify, profile and stratify populations by predicting future risk...allowing users to access a single, unified profile of an individual's historical care and predicted needs...to also forecast expenditure costs and medical service needs based on available data inputs." It helps professionals "determine which individuals are in need of specialized intervention programs and which intervention programs have the most impact on the quality of individuals' health."⁴⁰ ⁴¹ Optum serves more than 125M individual consumers, 9 out of 10 US hospitals, 4 out of 5 health plan organizations, and 33 State governments.⁴²

Although insurance companies and hospitals intend to align the interests of hospitals and patients, Obermeyer et al's research team surfaced a bias problem where at a given risk score, Black patients were considerably sicker than white patients with the same risk score. Across important markers of health needs—severity of diabetes, high blood pressure, renal failure, cholesterol, and anemia— Black patients were found to be substantially less healthy than white patients at any level of algorithm predictions. Impact

³⁶ Porter, Michael E., Thomas H. Lee (2013) The Strategy That Will Fix Health Care. Harvard Business Review. <https://hbr.org/2013/10/the-strategy-that-will-fix-health-care>

³⁷ Society of Actuaries (2016) Accuracy of Claims-Based Risk Scoring Models. <https://www.soa.org/globalassets/assets/files/research/research-2016-accuracy-claims-based-risk-scoring-models.pdf>

³⁸ Society of Actuaries (2016)

³⁹ Omale, Gloria (2019) The Need for AI Governance in Healthcare. Gartner. <https://www.gartner.com/smarterwithgartner/the-need-for-ai-governance-in-healthcare/>

⁴⁰ Impact Pro. Optum website. <https://www.optum.com/business/solutions/data-analytics/data-analytics-health-plans/impact-pro-cpl.html> (Accessed: August 16, 2021)

⁴¹ Impact Pro (2020) Better understand individual and population health needs. Optum. (Accessed: August 16, 2021) <https://www.optum.com/content/dam/optum3/optum/en/resources/sell-sheet/impact-pro-sell-sheet.pdf>

⁴² Optum website. <https://www.optum.com/about-us.html> (Accessed: August 16, 2021)

Pro relied on healthcare cost as a proxy for health and tried to take measures against racial discrimination by excluding race. Nevertheless, racial biases arose.⁴³ Research suggests that “[h]ealth care costs, though well measured and readily available in insurance claims data, are also the result of a complex aggregation process with a number of distortions due to structural inequality, incentives, and inefficiency. Effectively the risk scores were putting white patients higher in the queue for specialized interventions than black patients who were equally sick.”⁴⁴

In defining a problem, which algorithmic decision systems are being considered to solve, business leaders and their teams can begin by asking:

- Is this the correct definition of the problem?
- Who gets to define the problem and success criteria? How does this impact ownership of a problem?⁴⁵
- How will the framing of the problem impact the type of data we need to use? Will this introduce any biases into the system?
- What values are we embedding into the definition of the problem (both in development & implementation)?
- Who will benefit from this system and who might be harmed or might be at risk because of it?⁴⁶
- Do we know the historical, social, economic context of the issue at hand? How might this impact the data?⁴⁷
- Is this intended purpose clear? Can the system be used for other purposes?
- Are the incentive structures implicit in an algorithm? Can these contribute to unintentionally skewed outcomes?⁴⁸

For especially high-stakes situations, business leaders are encouraged to consider incorporating a process whereby the teams involved build out diagrams of the relationships between independent, control, and dependent variables (called “causal models” in

⁴³ Obermeyer et al (2019)

⁴⁴ Jee, Charlotte (2019) A biased medical algorithm favored white people for health-care programs. MIT Technology Review. <https://www.technologyreview.com/2019/10/25/132184/a-biased-medical-algorithm-favored-white-people-for-healthcare-programs>

⁴⁵ Elish, Madeleine Clare and Watkins, Elizabeth Anne (2020) *Repairing Innovation: A Study of Integrating AI in Clinical Care*. New York: Data & Society Research Institute. <https://datasociety.net/pubs/repairing-innovation.pdf>

⁴⁶ Cilauro, Federico, Montrasio, Francesca and Snelson, Sarah (2020) Algorithmic decision-making: Traps for the unwary. Frontier Economics. <https://www.frontier-economics.com/uk/en/news-and-articles/articles/article-i8417-empower-or-exclude/#>

⁴⁷ A more detailed section is provided at the end of this document for examples of structural and racial bias in the healthcare domain

⁴⁸ Cilauro, Federico, Montrasio, Francesca and Snelson, Sarah (2020)

research terms), involved in the system being developed. The development of these causal models should include robust discussion and documentation of two key features:

- the ***delayed impacts and feedback loops*** involved with complex, high-stakes contexts, and
- ***explicit references to any inferences*** being made regarding the system which allows for iterative learning to take place in partnership with outside stakeholders including policy makers and social groups most vulnerable to unfair outcomes⁴⁹.

Both developers and implementers of a system need to remember and incorporate these features into their processes. A system that was intended to demonstrate correlation(s) (variables that move in relation to one another as opposed to one causing another) might be wrongly used as if it was a causal model; a causal model might not take into account delayed impacts of certain important factors; or a causal model which is not transparent might be wrongly formulated to begin with or might drift in time due to continuous data inflow and learning. The use of the word “drift” in this sense means that the original intended use of the system drifts into other uses not intended when the system was built.

The Optum case is not the first, nor will it be the last example of unintentional harms caused by biased datasets and models created by humans. Though the researchers only focused on one particular tool, they identified the same flaw in the 10 most widely used algorithms in the industry.⁵⁰ The impact of this deficit is far-reaching: Each year, these tools are collectively applied to an estimated 150 to 200 million people, i.e., about half of the population, in the US.

⁴⁹ Martin, D., Prabhakaran, V., Kuhlberg, J.A., Smart, A., & Isaac, W.S. (2020) Participatory Problem Formulation for Fairer Machine Learning Through Community Based System Dynamics. ArXiv, abs/2005.07572.

⁵⁰ Obermeyer, Ziad. et al (2019)

Risks to the business

In the Fall of 2019, a series of articles in high profile sources such as the Washington Post⁵¹, Science⁵², Wall Street Journal⁵³, PBS,⁵⁴ and NBC⁵⁵ were published spotlighting the study published in Science. Optum responded that the study had “grossly mischaracterized” the situation based on incorrect use of the algorithm by one health system⁵⁶. This case study’s objective is to move the reader beyond the headlines and explain what went wrong across the lifecycle of a system and why it matters.

Better care, well-being, and dignity of the patients were some of the objectives of implementing this system, however, a functional risk might also arise due to the use of a system for different purposes. Especially in B2B (business-to-business) relationships, organizations are expected to have systems in place to monitor, audit, and improve their products, as well as ensure fidelity of use and outcomes with their clients. Identifying and removing these risks are the most crucial steps.

Any misalignment about the objectives can also magnify the reputational and financial risks. In terms of the reputation risk, once the popular press identifies a key case study, that case study is resurfaced repeatedly, cited again and again in the future news of biased algorithmic systems in the same domain, impacting the reputation of the company no matter how their representatives define the issue. Amazon’s AI recruiting tool that showed bias against women⁵⁷, for example, is continuously cited in the recruitment domain even though the company realized the issue itself, tried to correct it, and ultimately discontinued the tool.

UnitedHealth Group was also subjected to legal and regulatory risk. Following the publication of the research and publicity around it, Sens. Cory Booker, D-N.J., and Ron Wyden, D-Ore., published a series of letters asking The Centers for Medicare & Medicaid Services, the Federal Trade Commission, and the nation's five leading health insurers to

⁵¹ Johnson, Carolyn Y (2019) Racial bias in a medical algorithm favors white patients over sicker black patients. Washington Post.
<https://www.washingtonpost.com/health/2019/10/24/racial-bias-medical-algorithm-favors-white-patients-over-sicker-black-patients>

⁵² Price, Michael (2019) Hospital ‘risk scores’ prioritize white patients. Science.
<https://www.sciencemag.org/news/2019/10/hospital-risk-scores-prioritize-white-patients>

⁵³ Evans, Melanie, Mathews, Anna Wilde (2019) Researchers Find Racial Bias in Hospital Algorithm. WSJ.
<https://www.wsj.com/articles/researchers-find-racial-bias-in-hospital-algorithm-11571941096>

⁵⁴ Wu, Katherine J. (2019) Racially biased medical algorithm prioritizes white patients over black patients. PBS. <https://www.pbs.org/wgbh/nova/article/racial-bias-medical-algorithm-black-patients/>

⁵⁵ Gawronski, Quinn (2019) Racial bias found in widely used health care algorithm. NBC News.
<https://www.nbcnews.com/news/nbcblk/racial-bias-found-widely-used-health-care-algorithm-n1076436>

⁵⁶ McCullom, Rod (2020) Artificial Intelligence, Health Disparities, and Covid-19. Undark.
<https://undark.org/2020/07/27/ai-medicine-racial-bias-covid-19/>

⁵⁷ Dastin, Jeffrey (2018) Amazon scraps secret AI recruiting tool that showed bias against women. Reuters
<https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>

outline the actions they are taking to reduce algorithmic bias in healthcare.⁵⁸ Regulators started to review the details of this case in late 2019.

At a state level, Linda Lacewell, superintendent of the New York State Department of Financial Services, and Dr. Howard Zucker, State's health commissioner wrote to UnitedHealth Group chief executive David Wichmann,

"We call on you to immediately investigate these reports and demonstrate that this algorithm is not racially discriminatory or to cease using Impact Pro (or any other data analytics program) if you cannot demonstrate that it does not rely on racial biases or perpetuate racially disparate impacts."⁵⁹

Under state law, Optum could be liable for a fine of \$1,000 per incident if it was found to have breached anti-discrimination regulations. It is not possible to know the exact number of patients who could be impacted, however, even if only a fraction of Optum's 128 million customers were impacted, the financial burden would be significant.

There is also the impact to investors. Optum's recent announcement to acquire data analytics company Change Healthcare⁶⁰, according to some, would make them a "vertically integrated, all-seeing giant" in the industry by concentrating "an immense volume of competitively sensitive data" in their hands. This could sensitize investors to the importance of diligent handling of data as well as further increase regulatory scrutiny on the inner workings of Optum and how it will use this vast treasure trove of data to predict health risks.

Reputational, legal, and investor risks intersect. Public sentiment can drive regulatory investigations and that in turn can result in investor risk, all of which can lead to financial risk. As noted in legal risks, Optum faces the potential for significant fines from the state regulator as a direct result of this case. Indirectly, there may be fallout from customers (hospitals) in terms of a lack of trust in the product or the brand.

Optum's case involves considerable material risks, however, there may have been some external factors that tempered the extent to which the company has been adversely impacted to date. One factor was timing and the role of COVID-19, which diverted public attention as well as a regulatory focus away from the Optum case.

⁵⁸ Senate.gov (2019) Booker, Wyden Demand Answers on Biased Health Care Algorithms. <https://www.booker.senate.gov/news/press/booker-wyden-demand-answers-on-biased-health-care-algorithms>

⁵⁹ Snowbeck, Christopher (2019) NY Regulators Probe for Racial Bias in Health-Care Algorithm. Government Technology. <https://www.govtech.com/health/ny-regulators-probe-for-racial-bias-in-health-care-algorithm.html>

⁶⁰ Brown, Krista, Webb, Olivia (2021) The Big Tech of Health Care. The American Prospect. <https://prospect.org/health/big-tech-of-health-care-united-optum-change-merger/>

Another factor was Impact Pro's status as a supply chain component buried within the complexities of the healthcare system. Optum's customers are hospitals, not patients. There is an opaqueness to the situation in that patients would not know if they had been subjected to this particular algorithmic system, let alone whether they had been affected by this specific bias. Hospitals may have little incentive, outside of regulatory compliance, to change their operational systems. In this case, the bias in question came to light⁶¹ because the hospital in question was a large academic hospital that provided data to independent researchers for examination.

The authors believe it is currently premature to determine the final financial impact of the story on Optum. So far, the "fragmentation of American healthcare has kept Big Data from being fully harnessed"⁶² but it has also mitigated the damage that harmful/biased data applications can cause. Optum's acquisition of Change Healthcare leads to a significant increase in concentration and it could effectively create a new "Big Tech" of health care. After the merger, Optum/UnitedHealth will own not just the data, but also the providers, and the network through which people receive care. As power increases, so does the impact a single company can create as a benefit or damage, and thus the materiality of AI ethics. Optum/UnitedHealth's new status would be so powerful that it is hard to ignore the materiality of how it handles ethical issues and how it creates the boundaries and checks and balances between its different businesses.

Diversity of teams & stakeholders

A system that will serve and benefit the needs of diverse groups of people needs to incorporate the perspectives and experiences of those people and include their data in a representative way. A development approach, which does not take into account the impact of the system on those who are on the receiving end of the system's decisions, will be limited by the perspectives and values of its developers and funders. Therefore, it is crucial to map out accurately both internal and external stakeholders who could provide insights into the product development and ensure their participation. These stakeholders can help developers understand the context of data that is used and what might be the structural biases baked into the data, what might be the risks and harms if these are not corrected and provide feedback as to how the system is impacting different groups when it is in use.

Diversity in teams and stakeholders helps to challenge each other's assumptions regardless of whether those assumptions are conscious or unconscious. This check can help to mitigate risks further and avoid several consequences that might have gone

⁶¹ Jillson, Elisa (2021) Aiming for truth, fairness, and equity in your company's use of AI. Federal Trade Commission.
<https://www.ftc.gov/news-events/blogs/business-blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai>

⁶² Kotula, Henry (2021) Hospital mergers and acquisitions are a bad deal for patients. Why aren't they being stopped? <https://henrykotula.com/category/dept-of-justice-doj/>

unforeseen. As Ruha Benjamin, an Associate Professor of African American Studies at Princeton University, says of this algorithm “[w]hereas, in a previous era, the intention to deepen racial inequities was more explicit, today coded inequality is perpetuated precisely because those who design and adopt such tools are not thinking carefully about systemic racism.”⁶³

Optum states that they “routinely review and refine algorithms and tools to incorporate the latest research, data, and knowledge around their application. Teams of our researchers, physicians, and data scientists build and test new models, tools and designs not only to improve the ability to predict health needs but also to account for individual health variations and disparities to reduce bias.”⁶⁴ What other questions could be asked and answered by Optum in order to understand the diversity of opinions and experiences and stakeholder engagement involved with this system?

- How diverse is the development team?
- Who else is involved in the development process (for example, in this case study, who in addition to the above-mentioned experts were involved)?
- Does the organization have a stakeholder engagement policy?
- Does the organization engage and carefully consider diverse stakeholder perspectives in process/product design?
- Can those engaged in the development and deployment process freely flag concerns about risks and harms without whistleblowing?
- Who decides what feedback to incorporate into a product or service?

It is important that a system’s outcomes, including both the burdens and benefits, treat those impacted by those outcomes in a fair way. In this case, it should be expected that Impact Pro’s risk stratification tool provides outcomes about different patient groups in a non-discriminatory way. Where there is a genuine need for lawful bias, such bias should be transparent and explain publicly how it appears in the system and why it is justified. In fact, Optum acknowledges that “[i]t is critical to avoid gender, racial, and other forms of bias when using these types of algorithms [and that they] follow a rigorous process to create algorithms and health care metrics that are unbiased for their intended purpose.”⁶⁵ What other questions could be asked to ensure this result?

⁶³ Benjamin, Ruha (2019) Assessing risk, automating racism. Science.
<http://science.sciencemag.org/content/366/6464/421.abstract>

⁶⁴ Optum. Algorithms and population health management <https://www.optum.com/business/health-insights/algorithms-human-touch.html>

⁶⁵ Bgmanastiri (2019) Pro Case Study: Integrating claims and health risk assessment data to measure risk in a dual eligible population.
<https://bgmanastiri.info/content/dam/optum/resources/caseStudies/PRO%20Case%20Study.pdf>

- Is anyone or any group with subject matter expertise (history of gender or racial bias in the healthcare system, for example) involved in the design of the model?
- If there is input from a diverse set of stakeholders, are these valued and incorporated?
- What steps are taken following feedback from multi-stakeholder participation?
- Are these stakeholders given a chance to reassess the performance of products or services following implemented changes based on diverse inputs?
- How is the outcome quantified?

Within the health sector, using mortality or re-admission rates to measure hospital performance penalizes those serving poor or non-white populations.”⁶⁶ Unequal access to care and disparities across patients have complicated roots, including structural racism, income, education, lack of insurance, mistrust of the medical system, cultural misunderstandings, or even unconscious bias from the doctors’ side.⁶⁷

Incorporation of diversity and feedback requires working towards solutions and improvement of a system before, during and after it is commercially deployed. A harm might be unintended or unforeseen during a development stage. However, it is at the heart of responsible and ethical development and the use of algorithmic systems to mitigate bias and harm when they are made visible by different parties. It was only after a group of researchers completed their own independent analyses (as described above), that Optum was contacted by the researchers for an initial discussion of the results. In response, “the manufacturer [Optum/UnitedHealth Group] independently replicated analyses on its national dataset of 3,695,943 commercially insured patients’⁶⁸ ’ in order that the researchers and Optum could experiment with solutions together. Learning from the inappropriate use of a model and how past health care experiences *might* be presented in a biased manner that would point to critical questions which should be baked into whatever process changes Optum determines to make based on this case study.

⁶⁶ Obermeyer, Ziad. et al (2019)

⁶⁷ Johnson, Carolyn Y (2019)

⁶⁸ Obermeyer, Ziad. et al (2019)

Assurance & governance

Assurance of the quality of data, models, and systems as artifacts (i.e., anything that is created so a piece of software can be developed), requires transparency, demonstration of trustworthiness, and continuous improvement. AI assurance entails “a process that is applied at all stages of the AI engineering lifecycle ensuring that any intelligent system is producing outcomes that are valid, verified, data-driven, trustworthy and explainable to a layman, ethical in the context of its deployment, unbiased in its learning, and fair to its users.”⁶⁹

Diffusion of responsibility is a socio-psychological phenomenon where an individual abstains from taking action due to the belief that someone else is responsible for intervening.⁷⁰ Assurance is achieved by good governance which requires clear-cut responsibilities, accountability, and documentation. These need to be aligned with legal, regulatory, and compliance review, and the values of the organization. They need to be reflected across all relevant procedures and policies of the organization (like recruitment, compensation, and incentives, performance management, compliance, ability to raise concerns, etc.) and they need to be applied as business as usual, not as a stand-alone model only applied to data or algorithmic decision systems.

A paper entitled *Moral Crumple Zones*, written by cultural anthropologist Madeleine Clare Elish, advises that, as debates about the policy and ethical implications of AI systems grow, it will be increasingly important to accurately locate who is responsible when the agency is distributed in a system and control over an action is mediated through time and space. She describes how responsibility for an action may be misattributed to a human actor who had limited control over the behavior of an automated or autonomous system. The human in a highly complex and automated system may simply become a component—accidentally or intentionally—that bears the brunt of the moral and legal responsibilities when the overall system malfunctions. This protects the integrity of the technological system, at the expense of the nearest human operator. What is unique about the concept of a moral crumple zone is that it highlights how structural features of a system and the media’s portrayal of accidents may inadvertently take advantage of human operators (and their tendency to become “liability sponges”) to fill the gaps in accountability that may arise in the context of new and complex systems. The concept highlights how human operators may “absorb” responsibility when these gaps arise in socio-technical

⁶⁹ Batarseh, F.A., Freeman, L. & Huang, CH. (2021) A survey on artificial intelligence assurance. *J Big Data* 8, 60. <https://doi.org/10.1186/s40537-021-00445-7>

⁷⁰ Hooker, Sara (2021) Moving beyond “algorithmic bias is a data problem”. *Cell*. <https://doi.org/10.1016/j.patter.2021.10024>

systems in ways that do not reflect the distributed control and interactional aspects that compose the socio-technical system.⁷¹

The role of governance should not be just to assign liability. Governance should be understood as a system of responsibility-taking, authority, ownership, and values.⁷² Governance needs to fit the purpose of an AI system. It requires transparency and reporting, and it requires remediation when an artifact is found to be problematic. This approach goes beyond regulatory compliance. Just because a practice is not illegal does not mean that it is ethical or acceptable. Even if there is no regulatory requirement, a company takes steps beyond regulations to align its practices with its own policies and values. Governance applies to both those developing and implementing algorithmic systems. To answer some of these questions, specifically those that relate to data and models within similar context to Impact Pro, a team of researchers (some who were initially involved in the original Obermeyer et al research) are now developing a "playbook"⁷³ of best practices for auditing clinical algorithms, targeted for a variety of audiences -- including medical-algorithm designers, repository managers, manuscript writers and readers, journal editors, and model users. The long-term goal of the project is to develop a robust documentation standard. The playbook will include information and resources to help a healthcare system create and maintain its own "algorithm inventory" so it knows how and when software is used to make decisions and make ex ante and post hoc assessments. It will also cover how to monitor predictions made by the algorithm over time and across patient demographics, as well as how to assess an algorithm's performance based on what it is being used to predict or measure.⁷⁴ Knowing how the algorithmic tool is used can help organizations if the use is in alignment with any of the recommended uses of the tool.

⁷¹ Elish, Madeleine Clare (2019) Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction (pre-print) Engaging Science, Technology, and Society (pre-print), <http://dx.doi.org/10.2139/ssrn.2757236>

⁷² S&P Global (2020) What is the "G" in ESG? <https://www.spglobal.com/en/research-insights/articles/what-is-the-g-in-esg>

⁷³ Fenner, Marsha (2020) The MI-CLAIM checklist: Minimum information about clinical artificial intelligence modeling. Berkeley Institute for Data Science. <https://bids.berkeley.edu/news/mi-claim-checklist-minimum-information-about-clinical-artificial-intelligence-modeling>

⁷⁴ Ruiz, Rebecca (2020) Doctors use algorithms that aren't designed to treat all patients equally. Mashable. <https://mashable.com/article/healthcare-algorithms-racism/>

For any organization in an automated decision system ecosystem, it is therefore most crucial to first hire, train and develop internal talent who can ask the right questions, help build and execute the governance strategies and embed organization's policies and values into the processes. The importance of diversity of this talent across the lifecycle is already covered above. Here are some other key points with regards to internal talent:

- Are there clearly identified skills and/or experiences used in recruitment of those involved with algorithmic systems?
- Are the responsibilities and accountabilities clearly defined and communicated?
- Are developers and implementers incentivized (or reprimanded) to flag ethical concerns?
- Does the organization have a code of data ethics?
- How does the organization train its teams about bias, governance, assurance as well as keep current with new developments?

Once the crucial internal talent questions are addressed, some of the governance questions to ask are:

- What tools and information does the Board of Directors have to monitor these specific risks?
- Does the organization have an algorithmic risk committee?
- Where does the algorithmic risk committee (or roles/committees with similar responsibilities) sit in the organization chart? Is it empowered to make changes to mitigate or remove risks?
- Does the organization mandate impact and risk assessments for the models that it rolls out?
- How are these risks assessed, evaluated, monitored (ex-ante & post hoc)?
- How are the unintended unforeseen risks reviewed and mitigated? Who decides the mitigation plan?
- How does the organization document technical information about data and models and all the decisions made regarding the data, model architecture, and trade-offs?
- Does the organization report its results and impact transparently?
- Are the results and impact aligned to its ESG goals (if applicable)?
- Is there an ethical framework applied to the full lifecycle of the product?
- How does the organization proactively assess the socio-technical system in which its algorithm is deployed?
- How well does the model perform in different socioeconomic settings?
- Are there provisions to periodically reassess the performance of the product or services after preliminary review/roll out?
- Does the organization work with external third parties to audit its algorithms?
- How does the organization train the customers who use its products?
- How are shareholders and stakeholders (for example hospitals and patients) informed?
- How are shareholder and stakeholder concerns addressed and incorporated?

It is also crucial that the governance mechanisms in an organization cover data security, especially if there are no explicit data protection policies and laws in place⁷⁵. Even with a regulation like HIPAA (Health Insurance Portability and Accountability Act) in the US, it is often not clear what constitutes protected health information and what does not. This creates further complexities in the data share and quality check mechanisms. Electronic health record data are not consistently labeled, formatted, or routinely analyzable, “as they must be pulled and cleaned extensively from hospital data warehouses.”⁷⁶

⁷⁵ World Health Organization (2017) The health system and policy monitor: regulation. WHO. <https://www.who.int/health-laws/countries/usa-en.pdf?ua=1>

⁷⁶ Obermeyer, Ziad. et al (2019)

In almost all the domains where AI is becoming more prevalent, the regulation is still evolving. Just like in any other domain where emerging technologies (especially algorithmic systems) are impacting human lives and access to resources and opportunities, we need a combination of both ethical and responsible design and regulation with clear guidance, as minimum.

Data, processes & model details

As described earlier in this document, developers of the Impact Pro product used medical claims data to train the predictive models. Developers must be cognizant of the origins of the data, referred to as *data provenance*. Data provenance reflects the origination of the data at hand and the decisions, processes, inputs which altered the data throughout its lifecycle. When we decide to collect data, we make decisions on how to identify and classify the data points, where and how to collect data, what kind of consent is needed, what medium to use for collection, and what to include or exclude and how to process the collected data. If we are not collecting data ourselves but have opted for readily available datasets instead, we make decisions about which dataset to use, why that is a good dataset to answer our questions and who is represented in that dataset, and in which ways. We also take on all the data decisions and changes made by the original entity which created the dataset in the first place. In either case, we make decisions that reflect our experiences, knowledge and conscious, as well as unconscious, biases. Any dataset that reflects human behavior and interactions also reflect biases within the society that mirror the societal and historical imbalances.

Biases are encoded in historical healthcare data. Past patterns of care delivery can be captured in those algorithms and propagate previous biases. This can include the population upon which the model was trained. The representativeness of the data upon which the model was trained should be documented and described so that the limitations of the data (in terms of representativeness) are clearly known. That is, what is the balance of the different groups and how well does that replicate the proportion of the population that will be scored (that may include age, sex, race, ethnicity, and other socioeconomic features, or even geography). The balance of these groups defined by the labels must also be considered. For example, in the ImpactPro case, the developers might have considered the following:

- How are the outcomes / error rates different across different protected categories (False Positives, False Negatives)?
- Do patients with the same score have the same true need or outcome (irrespective of race or other social/political categories)?⁷⁷
- Should a single threshold (base rate) be used across different populations to achieve fairness?⁷⁸
- What is the prevalence of different groups in the data? How do we decide if it is sufficient?
- Is the representation across groups defined by other factors (age, geography, etc.) sufficient?⁷⁹
- How different is the population in the training dataset to the population in which the algorithm is being applied? Can this introduce bias into our outcomes?
- We often only discuss representativeness of data across one parameter (i.e., balance across sex). We must also consider intersectionality -- that is - - are Black women represented in the data? Or Latinx men, etc.?
- What is the prevalence of chronic kidney disease across groups defined by sex, race, and ethnicity?
- What is the prevalence of hypertension? Does it affect different groups at different or same/similar ages?
- What is the proportion of costs of care categorized by insurance status? How does this correlate to actual or anticipated income status of the insured? How does this vary across groups? How is the income status of the insured determined – self-declaration or estimated based on job title/position?
- Are there any possible biases in the model training, testing and validation processes?
- Is data provenance and changes to datasets clearly documented?

Additionally, they might have considered as the researchers have done:

- What is the severity of disease between Black and white individuals given the same risk score in the training data?

⁷⁷ Center for Applied AI at Chicago Booth. Algorithmic Bias Playbook.

https://www.ftc.gov/system/files/documents/public_events/1582978/algorithmic-bias-playbook.pdf

⁷⁸ Kleinberg, J., Mullainathan, S., & Raghavan, M. (2017). Inherent Trade-Offs in the Fair Determination of Risk Scores. *ArXiv, abs/1609.05807*

⁷⁹ Another recent study published in *JAMA* (the *Journal of the American Medical Association*), over 70 publications were reviewed focused on algorithms used across several areas of clinical medicine. Most of the data used to train those AI algorithms came from just three states: California, New York, and Massachusetts.

Descriptive analyses such as these can assist the developers in understanding the imbalance in their data and highlight the concerns that can emerge once a model is in production.

Lack of diversity and representation of marginalized groups, across geographical regions, is endemic in datasets in almost all algorithmic applications such as facial recognition, natural language processing, and image recognition.

In ImpactPro's case, race was deliberately omitted as a predictive feature. However, the omission of race does not guarantee model fairness across groups of individuals defined by race, as becomes evident in this case study (i.e., fairness through unawareness cannot be assumed). In order to consider fairness across model performance, there are many protected features that developers are cognizant of, such as sex, race, ethnicity, and age (depending on the use case). Many surrogate (proxy) variables must also be considered - those that are closely correlated to and may even be predictive of protected features. These may include features that encode an individual's geography such as zip code or census tract and those that indicate access to care such as insurance status and prior medical costs.

There is much debate about the use of protected features such as race and ethnicity as variables in predictive models. While the full discussion is beyond the scope of this paper, it may suffice to say that a thoughtful discussion with representative stakeholders is recommended before this determination is made. We must go beyond data science to ask the hard questions and to determine who will be served and harmed by the decisions made. Several guidelines delve deeper into bias mitigation in datasets.^{80 81 82} This is as important in health care as it is in housing, employment, financial services, and criminal justice, all of which also use predictive algorithms as decision-making tools.

Deployment

According to Optum's website, the Impact Pro product "can determine which individuals are in need of specialized intervention programs and which intervention programs are likely to have an impact on the quality of individuals' health."⁸³ According to Tyler Mason, UnitedHealth Group's VP of Communications, the Impact Pro tool was "to be used only to determine potential future health care costs for individual patients based on past health

⁸⁰ Brown, S., Carrier, R., Hickok, M., & Smith, A. L. (2021). Bias Mitigation in Data Sets. <https://doi.org/10.31235/osf.io/z8qrb>

⁸¹ Alethicist.org. Ethical AI Frameworks, Guidelines, Toolkits Repository. <https://www.aiethicist.org/frameworks-guidelines-toolkits>

⁸² Berkeley Haas Center for Equity, Gender, and Leadership (EGAL) Equity Fluent Leadership (EFL) Playbook: Mitigating Bias in Artificial Intelligence (AI) <https://haas.berkeley.edu/equity/industry/playbooks/mitigating-bias-in-ai/>

⁸³ Optum ImpactPro. <https://www.optum.com/business/solutions/health-plans/data-analytics/impact-pro.html>

care experiences and that it was not intended to assist doctors in making care decisions.”⁸⁴ A page in Optum’s website suggests they “train and strongly counsel those using our tools to use them as advised and to take into account a range of information about a patient - such as demographics, health care utilization, social determinants of health, and gaps in care measures, as well as expert observations of human clinicians of the patient - when making care decisions.”⁸⁵

To ensure intended usage of an algorithmic tool, we recommend a clear communication between the developing and implementing organizations and encourage them to become stakeholders in each other’s processes.

As experts from around the world gather to determine principles necessary to guide behavior and inform standards and policy making, a few common imperatives have been raised repetitively. One of those imperatives is the concept of *effectiveness* and that the creators and operators of AI and intelligent systems provide evidence of the effectiveness and fitness for purpose of the models they create.⁸⁶ This imperative could be applied by asking the following questions:

- What metrics and or benchmarks will serve as gauges of the effectiveness of the system?
- Did the developers contemplate that the intended use of the tool would/could “drift” (for example, that health care user objectives and circumstances would be different than the intended use of the model)?
- How does the organization train the implementers of its systems and ensure they understand the risks and potential shortcomings of the system and intended use?
- What human review and oversight mechanisms are in place?
- Does the organization / customer have the ability to stop the system if any safety issue is identified?
- Did the developers of the system provide guidance on how to interpret and respond to the metrics generated by the systems?⁸⁷

⁸⁴ McCullom, Rod (2020) <https://undark.org/2020/07/27/ai-medicine-racial-bias-covid-19/>

⁸⁵ Optum. Algorithms and population health management <https://www.optum.com/business/health-insights/algorithms-human-touch.html> (Note: date the page is created is not available. Last accessed on August 16th, 2021)

⁸⁶ The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems (2019) Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems, First Edition. IEEE. <https://standards.ieee.org/content/ieee-standards/en/industry-connections/ec/autonomous-systems.html>

⁸⁷ IEEE (2019), EAD1e

Obermeyer et al advised that Optum was responsive to their research and an Optum team worked with the researchers to improve the model.⁸⁸ For the purposes of this case study, one can ask how can patients be assured now that Optum has communicated with all the hospitals using its model and rolled out the improved version(s)? Some additional questions can help organizations to ensure the tools they developed are deployed correctly.

- How are the customers using the tools the organization developed?
- After deployment, what mechanisms should be in place to ensure customers' use of the tool is within recommended purposes?
- What guidance should the organization offer to its customers?
- What is the risk to the organization if a customer operates the tool outside of its recommended purposes?
- What is the organization's responsibility to those ultimately impacted if the tool is used outside of its recommended purposes?

In a draft proposal to identify and manage bias in AI dated June 2021, the United States' National Institute for Standards and Technology ("NIST") discusses how the deployment stage offers us "an interesting window into how perceptions and uses can differ...once the technology is deployed and used in different settings and for different purposes, we see perception turn to unintended use causes and even distrust."⁸⁹ NIST further points out:

"[a]nother important gap that contributes to bias relates to differences in interpretability requirements between users and developers...There are individual differences in how humans interpret AI model output. When system designers do not take these differences into consideration it can contribute to misinterpretation of that output.⁹⁰ When these differences are combined with the societal biases found in datasets and human cognitive biases such as automation complacency (which is particularly relevant in the deployment stage), where end users may unintentionally 'offload' their decisions to the automated tool - this can cause significant negative impacts."⁹¹

⁸⁸ CogX Festival 2021. Session: "Ethics: Getting to 'good outcomes' from AI-generated decisions: What are they and who decides?" <https://www.youtube.com/watch?v=2OjtgsI4IOI>

⁸⁹ Schwartz, Reva, Down, Leann, Jonas, Adam, Tabassi, Elham (2021) Draft NIST Special Publication 1270 A Proposal for Identifying and Managing Bias in Artificial Intelligence <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270-draft.pdf>

⁹⁰ Broniatowski, D. (2021), Psychological Foundations of Explainability and Interpretability in Artificial Intelligence, NIST Interagency/Internal Report (NISTIR), National Institute of Standards and Technology, Gaithersburg, MD, [online], <https://doi.org/10.6028/NIST.IR.8367>

⁹¹ Schwartz et al (2021)

Monitoring, feedback, calibration / improvement

Developing and deploying an algorithmic decision tool is only one part of the lifecycle of a product. Once a tool is put into use, the post-market product must be monitored for performance as per agreed-upon metrics. Feedback should be reviewed, and where necessary, the model should be calibrated or improved upon to ensure that it still meets the purpose and that emerging risks are mitigated. These actions are still part of the overall governance process. However, special attention should be paid here to **feedback**. The concept can have two meanings in the context of automated decision systems.

The first meaning relates to the concerns, issues or even positive comments emerging from those using a system or being impacted by it, as well as the mechanism to make sense of these comments. After a model is live and performing, a previously unidentified consequence might surface. It is crucial for system owners (vendor & customer) to ensure that there are mechanisms to collect this information and correct and/or improve a model's outcomes, efficiency, and effectiveness.

It is the responsibility of vendors to ensure the users are trained in the best way when systems are deployed, that the users are aware of the purpose of the system, the design, the decisions, and shortcomings of the system. We know from Obermeyer et al's statements that "the manufacturer [Optum] independently replicated analyses on its national dataset...To resolve the issue, [researchers & Optum] began to experiment with solutions together."

In this phase, business leaders and their teams can ask:

- How is the performance of the product or service evaluated?
- How are the criteria for such evaluation determined?
- What feedback and communication systems are in place between the organization and customer to allow for continuous feedback & improvement?
- What monitoring mechanisms are in place to ensure the tool's outcomes meet the agreed-upon performance metrics?
- What is the protocol to pause/discontinue the product/service when concerns are raised?

The second meaning of feedback as it relates to automated decision systems is to view it as a data input-output loop: For example, as data is inputted into a system to create outputs, the outputs themselves might then be fed back into the system as inputs so whatever is captured within the dataset replicates itself. This can magnify or deepen the biases and risks along with it.

Alternatively, if a machine learning system is being used, the algorithm learns from user behavior and feeds that behavior back into the model. For example, the system only

recommends the same types of books, movies, news, jobs a person was interested in before. It might learn the behavior of a clinical staff (or a group of staff in a clinic for example) and only suggest similar results in the future.

Algorithm performance can also change when the underlying data changes (for example, when a hospital starts using new medical imaging technology) or when algorithms are used in new locations and/or in different contexts (for example, pediatric vs. geriatric or inpatient vs outpatient populations).

- Does the team monitoring the performance of the algorithmic tool have any predefined set of metrics to trigger a review of drift?⁹²

To be successful, AI interventions must always be thought of as sociotechnical systems, in which social context, relationships, and power dynamics are central, not an afterthought.⁹³ Datasets used in these tools reflect how individuals interact with each other and with institutions in ways that are shaped by previous experiences. Algorithms, with their scale and efficiency in processing, amplify these biases. They create feedback loops that not only widen the divides in the society, but also bring about new and more complicated ways individuals and underserved groups might be discriminated against.

In the use of Impact Pro in this case, differences between various groups in terms of their access (or ability to access) health systems and the historical interactions of Black patients with health establishment should have been taken into account. Health care costs, mainly reflected in insurance claims data, are also the result of a complex aggregation process with a number of distortions due to structural inequality, incentives, and inefficiency.⁹⁴

Although race was excluded as a category in the model, the bias surfaced in relation to race. It was failing to account for nearly 50,000 chronic conditions experienced by Black patients. They found substantial disparities in health conditional on risk but little disparity in costs. In other words, the model used by the healthcare system was great at predicting cost but not health needs. However, other factors could have also surfaced. Socio-economic conditions in terms of income level, access to insurance and transportation, or childcare support can lead to differences in healthcare.

When looking at feedback as a data input-output loop, different factors, their correlation, and interaction with the outcomes should also be analyzed. As the datasets and the conditions change, it is also critical to assess at regular intervals how best the model adjusts to such changes.

⁹² Center for Applied AI at Chicago Booth. Algorithmic Bias Playbook

⁹³ Elish, Madeleine Clare and Watkins, Elizabeth Anne (2020) *Repairing Innovation: A Study of Integrating AI in Clinical Care* (New York: Data & Society Research Institute). <https://datasociety.net/pubs/repairing-innovation.pdf>

⁹⁴ Obermeyer, Ziad. et al (2019)

Know the context of your data historical / structural bias in healthcare

We mentioned previously that data provenance reflects the origination of the data at hand and the decisions, processes, inputs which altered the data throughout its lifecycle. It is crucial for data scientists, algorithmic tool developers and implementers to understand the context of data and the historical and structural problems that might be replicated.⁹⁵ Knowing the context of the data allows designers and decision-makers to frame the question(s) correctly, have a better understanding of who should be involved in the process, identify possible biases in the data, and learn from the lessons of failures and successes of the past. Without knowing the context of data and the problem, not only will the solutions not be optimal, but the necessary safeguards to protect individuals, groups or societies will not be in place. An organization which does not take into consideration the history of the problem only exposes itself to more risks.

In 2017 UnitedHealth (Optum's parent company) was one of the subjects of another case study "The Competitive Advantage of Racial Equity"⁹⁶ and a 2019 follow up report "How Advancing Racial Equity Can Create Business Value"⁹⁷ which confirms that the organization recognizes that health outcomes are not only influenced by clinical factors but also by social determinants that include race and the overall effects of structural racism,⁹⁸ and by the admission of Michael Currie, SVP and Chief Health Equity Officer of United Health Group "It's never a matter of if health disparities exist, but where and at what magnitude they exist."⁹⁹

To provide some further context for those readers who might be interested, "there is a long history of Black patients facing barriers to accessing equitable care¹⁰⁰ receiving less desirable treatments or race becoming a direct factor in healthcare decisions.¹⁰¹ Structural racism, education, mistrust of the medical system, segregated hospital facilities, racist

⁹⁵ Thesmar, D., Sraer, D., Pinheiro, L. *et al.* (2019) Combining the Power of Artificial Intelligence with the Richness of Healthcare Claims Data: Opportunities and Challenges. *PharmacoEconomics* 37. <https://doi.org/10.1007/s40273-019-00777-6>

⁹⁶ Blackwell, Angela Glover, Kramer, Mark, Vaidyanathan, Lalitha, Iyer, Lakshmi, Kirschenbaum, Josh (2019) Health Care and the Competitive Advantage of Racial Equity <https://www.fsg.org/publications/competitive-advantage-racial-equity>

⁹⁷ *Id.*

⁹⁸ *Id.*

⁹⁹ *Id.*

¹⁰⁰ Johnson, Carolyn Y (2015) Racial inequality even affects how long we wait for the doctor. Washington Post. https://www.washingtonpost.com/news/wonk/wp/2015/10/05/medicines-racial-divide-measured-in-minutes-spent-waiting-to-see-a-doctor/?itid=lk_inline_manual_13

¹⁰¹ Malcom, Kelly (2020) Racially Biased Oxygen Readings Could Be Putting Patients at Risk. University of Michigan Health Lab. <https://labblog.uofmhealth.org/rounds/racially-biased-oxygen-readings-could-be-putting-patients-at-risk>

medical curricula,¹⁰² cultural misunderstandings, or even unconscious biases¹⁰³ that doctors may not even know they have¹⁰⁴ can also contribute to discrepancies in care.¹⁰⁵ Collectively, this means that healthcare spending for Black patients is significantly less even though they might have the same level of need. Studies have found Black patients are less likely to receive pain¹⁰⁶ treatment, potentially lifesaving surgery¹⁰⁷ or cholesterol-lowering drugs, compared with white patients. Khiara M. Bridges, professor of law and professor of anthropology at Boston University, summarizes research by the Institute of Medicine (now called National Academy of Medicine (NAM)) that found “racial and ethnic minorities receive lower-quality health care than white people—even when insurance status, income, age, and severity of conditions are comparable.” By “lower-quality health care,” NAM meant the concrete, inferior care that physicians give their Black patients with regards to appropriate cardiac care, kidney dialysis or transplants and best treatments for stroke, cancer, or AIDS.¹⁰⁸ Another study of 400 hospitals in the United States showed that Black patients with heart disease received older, cheaper, and more conservative treatments than their white counterparts¹⁰⁹. Black patients were less likely to receive coronary bypass operations and angiography; that after surgery, they are discharged earlier from the hospital than white patients—at a stage when discharge is inappropriate; that Black patients are more likely to receive fewer desirable treatments. The rates at which Black patients have their limbs amputated is higher than those for white patients.¹¹⁰ A systematic review of pain management studies found that Black patients were 40

¹⁰² Benjamin, Ruha (2019)

¹⁰³ Hall WJ, Chapman MV, Lee KM, et al. (2015) Implicit Racial/Ethnic Bias Among Health Care Professionals and Its Influence on Health Care Outcomes: A Systematic Review. *Am J Public Health*. 105(12):e60-e76. doi:10.2105/AJPH.2015.302903

¹⁰⁴ Johnson, Carolyn Y (2019)

¹⁰⁵ Id.

¹⁰⁶ Somashekhar, Sandhya (2016) The disturbing reason some African American patients may be undertreated for pain. Washington Post. https://www.washingtonpost.com/news/to-your-health/wp/2016/04/04/do-blacks-feel-less-pain-than-whites-their-doctors-may-think-so/?itid=lk_inline_manual_13

¹⁰⁷ Bach, Peter B.Cramer, Laura D.Warren, Joan L.Begg, Colin B. (1999) Racial Differences in the Treatment of Early-Stage Lung Cancer. *N Engl J Med* 1999; 341:1198-1205. Massachusetts Medical Society. <https://doi.org/10.1056/NEJM199910143411606>

¹⁰⁸ Bridges, Khiara M. (2018) Implicit Bias and Racial Disparities in Health Care. *Human Rights Magazine* Vol. 43, No. 3: The State of Healthcare in the United States. https://www.americanbar.org/groups/crsj/publications/human_rights_magazine_home/the-state-of-healthcare-in-the-united-states/racial-disparities-in-health-care/

¹⁰⁹ Breathett K, Yee E, Pool N, et al. (2019) Does Race Influence Decision Making for Advanced Heart Failure Therapies? *J Am Heart Assoc.* ;8(22):e013592. doi:10.1161/JAHA.119.013592

¹¹⁰ Bridges, Khiara M. (2018)

percent less likely to receive pain medication in U.S. emergency departments as compared to white patients¹¹¹. Hispanic patients were 25 percent less likely.¹¹²

Another way race is inserted into health and medicine is when diagnostic algorithms and practice guidelines adjust or “correct”¹¹³ their outputs based on a patient’s race or ethnicity.¹¹⁴ Many of these race-adjusted algorithms guide decisions in ways that may direct more attention or resources to white patients than to members of racial and ethnic minorities. Race adjustments that yield higher estimates of kidney function in Black patients might delay their referral for specialist care or transplantation and lead to worse outcomes. The Vaginal Birth after Cesarean (VBAC) algorithm predicts the risk posed by a trial of labor for someone who has previously undergone cesarean section. It predicts a lower likelihood of success for anyone identified as African American or Hispanic.¹¹⁵

The COVID-19 pandemic has brought social and racial injustice and inequity to the forefront of public health. It has highlighted that health equity is still not a reality as COVID-19 has unequally affected many racial and ethnic minority groups^{116 117 118}, putting them more at risk of getting sick and dying from the virus. In the coming years, the data generated during this period may significantly impact the results in unexpected ways. We might experience issues and disparities in data similar to the ones analyzed in this case study due to COVID-19 pandemic. Some questions we could be asking now include:

¹¹¹ Lee, Paulyne, Le Saux, Maxine, Siegel, Rebecca, Goyal, Monika, Chen, Chen, Ma, Yan, Meltzer, Andrew C. (2019) Racial and ethnic disparities in the management of acute pain in US emergency departments: Meta-analysis and systematic review. *The American Journal of Emergency Medicine* Vol.37-9. <https://doi.org/10.1016/j.ajem.2019.06.014>

¹¹² Watson, Wendy, Marsh, Christina () Artificial Intelligence Bias in Healthcare. Booz Allen Hamilton Perspectives.

<https://www.boozallen.com/c/insight/blog/ai-bias-in-healthcare.html>

¹¹³ Madhusoodanan, Jyoti (2020) Is a racially biased algorithm delaying health care for one million Black people? *Nature* 588, 546-547 <https://doi.org/10.1038/d41586-020-03419-6>

¹¹⁴ Vyas, Darshali A., Eisenstein, Leo G., Jones, David S. (2020) Hidden in Plain Sight — Reconsidering the Use of Race Correction in Clinical Algorithms *New England Journal of Medicine* Vol 383-9. <https://www.nejm.org/doi/full/10.1056/NEJMms2004740>

¹¹⁵ Waddell, Kaveh (2020) Medical Algorithms Have a Race Problem. *Consumer Reports* <https://www.consumerreports.org/medical-tests/medical-algorithms-have-a-race-problem/>

¹¹⁶ Centers for Disease Control and Prevention (2019) Health Equity Considerations and Racial and Ethnic Minority Groups. <https://www.cdc.gov/coronavirus/2019-ncov/community/health-equity/race-ethnicity.html#fn2>

¹¹⁷ The COVID Racial Data Tracker. The Atlantic. <https://covidtracking.com/race>

¹¹⁸ Chin-Hong, P., Alexander, K.M., Haynes, N. *et al.* (2020) Pulling at the heart: COVID-19, race/ethnicity, and ongoing disparities. *Nat Rev Cardiol* 17. <https://doi.org/10.1038/s41569-020-0416-6>

- Can data collected during COVID-19 treatment impact results / insurance coverage (access to co-morbidity data)?
- Can imbalanced health access of certain groups during COVID-19 amplify the bias in data?¹¹⁹
- How does data representation change when adjusted for different race or age groups? Or when data is (dis)aggregated?¹²⁰
- Can COVID-19 change the way the data is interpreted, or the way issues of bias are discussed?
- Have risk predictions contributed to the racial disparities seen in Covid-19 treatment and access to testing?¹²¹
- How might the decisions around where and how to collect data lead to defining what was collected (from issues of transportation, bias of drive-thru testing, where sites were established etc.).^{122 123}
- How will telehealth data be integrated into health access data?¹²⁴
- What are / were the implications of COVID-19 for social determinants of health?¹²⁵

Conclusion

As evidenced by the examples of bias at the beginning of this paper, the issues covered here are not only specific to the health domain, or that Optum's software is the only case of a system resulting in biased decisions. This case study can be applied to different systems across different industries like employment, credit scoring, housing, access to education, criminal justice, and more. As algorithmic decision tools, like risk scoring or resource allocation, are becoming more ubiquitous in both public and private spheres, the differential treatment of humans and possible discrimination due to age, race, ethnic

¹¹⁹ Eliane Röösl, Brian Rice, Tina Hernandez-Boussard, Bias at warp speed: how AI may contribute to the disparities gap in the time of COVID-19, *Journal of the American Medical Informatics Association*, Volume 28, Issue 1, January 2021, Pages 190–192, <https://doi.org/10.1093/jamia/ocaa210>

¹²⁰ APM Research Lab (2021) The color of coronavirus: COVID-19 deaths by race and ethnicity in the U.S. <https://www.apmresearchlab.org/covid/deaths-by-race>

¹²¹ McCullom, Rod (2020) Is artificial intelligence worsening COVID-19's toll on Black Americans? Massive Science <https://massivesci.com/articles/ai-medicine-racial-bias-covid-19/>

¹²² Adeline, Stephanie (2020) In Large Texas Cities, Access To Coronavirus Testing May Depend On Where You Live. NPR. <https://www.npr.org/sections/health-shots/2020/05/27/862215848/across-texas-black-and-hispanic-neighborhoods-have-fewer-coronavirus-testing-sit>

¹²³ McCullom, Rod (2020) <https://massivesci.com/articles/ai-medicine-racial-bias-covid-19/>

¹²⁴ Wicklund, Eric (2021) With ER Visits Rising, Payers Pitch mHealth and Telehealth as an Alternative. MHealth Intelligence. <https://mhealthintelligence.com/news/with-er-visits-rising-payers-pitch-mhealth-and-telehealth-as-an-alternative>

¹²⁵ Orgera, Kendal, Garfield, Rachel, Rudowitz, Robin (2021) Implications of COVID-19 for Social Determinants of Health. Kaiser Family Foundation. <https://www.kff.org/coronavirus-covid-19/issue-brief/implications-of-covid-19-for-social-determinants-of-health/>

background, sex, gender identity, ability, citizenship status (just to name a few) will be more and more automated.

Optum's case is neither the first nor will it be the last. Not even in the healthcare domain. In fact, while our team was writing this case study, a new study was published to advise that Epic Systems' alert system could only correctly differentiate patients who do and don't have sepsis 63 percent of the time and that part of the problem seemed to be in the way the Epic algorithm was developed. It defined sepsis based on when a doctor would submit a bill for treatment, not necessarily when a patient first developed symptoms.¹²⁶¹²⁷ Epic Systems is one of the biggest electronic health record companies in the United States. An Epic spokesperson disputed the findings¹²⁸ saying that another research showed the algorithm was accurate.

Ruha Benjamin puts it eloquently and succinctly: "This study contributes greatly to a more socially conscious approach to technology development, demonstrating how a seemingly benign choice of label (that is, health cost) initiates a process with potentially life-threatening results...[D]epending on how one defines the problem to be solved—whether to lower health care costs or to increase access to care—the outcomes will vary considerably. [This means] it is vital to develop tools that move from assessing individual risk to evaluating the production of risk by institutions so that, ultimately, the public can hold them accountable for harmful outcomes."¹²⁹

Ethical and responsible development and implementation of technologies requires value alignment, deep effort to understand the context and the problem, intentional steps to mitigate and remove risks, and robust governance and accountability mechanisms. Internal and external validation are necessary to be sure that decisions are being made with the highest quality evidence.¹³⁰ Monitoring and evaluation strategies should be in place that subscribe to latest practices within and outside the domain of application to ensure dataset quality and diversity. Ethics is not a nice to have. It is a strategic decision for the business. It requires vision. It is a responsibility towards society and the future - the future generations, the future planet, and the future we imagine.

¹²⁶ Wetsman, Nicole (2021) A hospital algorithm designed to predict a deadly condition misses most cases. The Verge.

<https://www.theverge.com/2021/6/22/22545044/algorithm-hospital-sepsis-epic-prediction>

¹²⁷ Wong A, Otles E, Donnelly JP, et al.(2021) External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. *JAMA Intern Med.*;181(8):1065–1070. doi:10.1001/jamainternmed.2021.2626

¹²⁸ Ross, Casey (2021) A popular algorithm to predict sepsis misses most cases and sends frequent false alarms, study finds. StatNews.<https://www.statnews.com/2021/06/21/epic-sepsis-prediction-tool/>

¹²⁹ Benjamin, Ruha (2019)

¹³⁰ Brnabic, A., Hess, L.M. (2021) Systematic literature review of machine learning methods used in the analysis of real-world data for patient-provider decision making. *BMC Med Inform Decis Mak* 21, 54. <https://doi.org/10.1186/s12911-021-01403-2>

Information on Optum

Optum is a health services business serving the broad healthcare marketplace, including payers, care providers, employers, governments, life sciences companies and consumers, through its OptumHealth, OptumInsight and OptumRx businesses¹³¹. These businesses have dedicated units that help improve overall health system performance through optimizing care quality, reducing costs and improving consumer experience and care provider performance, leveraging distinctive capabilities in data and analytics, pharmacy care services, population health, health care delivery and health care operations

UnitedHealthcare and Optum, in 2019, processed nearly a trillion dollars in gross billed charges and managed more than \$250 billion in aggregate health care spending on behalf of the customers and consumers. Revenues are derived from premiums on risk-based products; fees from management, administrative, technology, consulting, and managed outsourced services; sales of a wide variety of products and services related to the broad health care industry; and investment and other income.

Optum operates three business segments leveraging distinctive capabilities in data and analytics, pharmacy care services, population health, health care delivery and health care operations:

- OptumHealth focuses on care delivery, care management, wellness and consumer engagement, and health financial services (serving 9 out of 10 U.S. hospitals and more than 100,000 physicians)
- OptumInsight offers data, analytics, research, consulting, technology, and managed services solutions (serving 4 out of 5 U.S. health plans to improve efficiency, understand and optimize growth while managing risk, improve payment integrity.
- OptumRx provides a diversified array of pharmacy care services.

=====

Acknowledgements

Authors would like to thank Patricia Shaw and Cortnie Abercrombie for their review and support.

¹³¹ UnitedHealth Group (UNH) Annual Reports 2019. <https://www.unitedhealthgroup.com/investors/annual-reports.html>