

# Keyon Vafa

---

Postdoctoral Fellow  
Data Science Initiative  
Harvard University

[kvafa@g.harvard.edu](mailto:kvafa@g.harvard.edu)  
[www.keyonvafa.com](http://www.keyonvafa.com)

Education	<b>Columbia University</b> Ph.D. Computer Science Thesis: Interpretable Machine Learning for the Social Sciences: Applications in Political Science and Labor Economics Committee: David Blei (advisor), Susan Athey, Suresh Naidu, Zhou Yu, Richard Zemel	2017 - 2023
	<b>Columbia University</b> M.S. Computer Science	2017 - 2018
	<b>Harvard University</b> B.A. Computer Science and Statistics, <i>magna cum laude</i>	2012 - 2016
Awards and Fellowships	Harvard Data Science Initiative Postdoctoral Fellowship	2023 -
	Cheung-Kong Innovation Doctoral Fellowship	2020 - 2022
	Columbia University Nominee for Google PhD Fellowship	2019
	National Science Foundation, Graduate Research Fellowship	2016 - 2019
	Phi Beta Kappa Society	2016
	Bok Center Certificate of Distinction in Teaching	2015
	John Harvard Scholar (grade point average in top 5% of class)	2013 - 2015
Selected Papers	<b>K. Vafa</b> . Is causal inference compatible with frictionless reproducibility? <i>Harvard Data Science Review</i> [to appear].	2024
	<b>K. Vafa</b> , S. Athey, D. Blei. <a href="#">Decomposing changes in the gender wage gap over worker careers</a> . <i>National Bureau of Economic Research</i> .	2023
	<b>K. Vafa</b> , E. Palikot, T. Du, A. Kanodia, S. Athey, D. Blei. <a href="#">CAREER: A foundation model for labor sequence data</a> . <i>Transactions of Machine Learning Research (TMLR)</i> [previously spotlight presentation at NeurIPS Workshop on Distribution Shifts].	2023
	C. Zheng, <b>K. Vafa</b> , D. Blei. <a href="#">Revisiting topic-guided language models</a> . <i>Transactions of Machine Learning Research (TMLR)</i> .	2023
	C. Zheng, C. Shi, <b>K. Vafa</b> , A. Feder, D. Blei. <a href="#">An invariant learning characterization of controlled text generation</a> . <i>Proceedings of the Association for Computational Linguistics (ACL)</i> .	2023

- K. Vafa.** [Interpretable machine learning for the social sciences: Applications in political science and labor economics.](#) *Ph.D. Thesis.* 2023
- K. Vafa,** Y. Deng, D. Blei, A. Rush. [Rationales for sequential predictions.](#) 2021  
*Proceedings of Empirical Methods in Natural Language Processing (EMNLP).*
- A. Schein, **K. Vafa,** D. Sridhar, V. Veitch, J. Moffet, J. Quinn, N. Makiya, D. Blei, Donald Green. [A digital field experiment reveals large effects of friend-to-friend texting on voter turnout.](#) *The Web Conference (WWW).* 2021
- K. Vafa,** S. Naidu, D. Blei. [Text-based ideal points.](#) *Proceedings of the Association for Computational Linguistics (ACL).* 2020
- D. Tran, **K. Vafa,** K. Agrawal, L. Dinh, B. Poole. [Discrete flows: Invertible generative models of discrete data.](#) *Proceedings of Neural Information Processing Systems (NeurIPS).* 2019
- K. Vafa.** [Training deep Gaussian processes with sampling.](#) *NeurIPS Workshop on Advances in Approximate Bayesian Inference Workshop.* 2016
- K. Vafa,** C. Haigh, A. Leung, N. Yonack. [Price discrimination in the Princeton Review's online SAT tutoring service.](#) *Journal of Technology Science.* 2015

## Selected Talks

- [Decomposing Changes in the Gender Wage Gap over Worker Careers,](#) 2023  
[Workshop in AI + Economics in Zurich](#)
- [Decomposing the Gender Wage Gap with a Foundation Model of Labor History,](#) 2023  
[Copenhagen Center for Social Data Science at the University of Copenhagen](#)
- [Decomposing the Gender Wage Gap with a Foundation Model of Labor History,](#) 2023  
[Technical University of Denmark \(DTU\)](#)
- [Decomposing Changes in the Gender Wage Gap over Worker Careers,](#) 2023  
[Harvard IQSS Workshop in Applied Statistics](#)
- [Foundation Models for Labor Economics,](#) 2023  
[Machine Learning in Economics Summer Institute at The University of Chicago Booth](#)
- [Decomposing Changes in the Gender Wage Gap over Worker Careers,](#) 2023  
[NBER Summer Institute 2023 \(Labor Studies\)](#)
- [Adjusting the Gender Wage Gap with a Low-Dimensional Representation of Job History,](#) 2023  
[Deepmind](#)
- [Adjusting the Gender Wage Gap with a Low-Dimensional Representation of Job History,](#) 2023  
[Stanford University](#)

CAREER: Transfer Learning for Economic Prediction of Labor Sequence Data, University of Chicago	2023
CAREER: Transfer Learning for Economic Prediction of Labor Sequence Data, <a href="#">ETH Zurich AI Center</a>	2023
CAREER: Transfer Learning for Economic Prediction of Labor Sequence Data, <a href="#">Microsoft Research Computational Social Science group</a>	2023
CAREER: Economic Prediction of Labor Sequence Data Under Distribution Shift (spotlight talk), <a href="#">NeurIPS Workshop on Distribution Shifts</a>	2022
Learning a Low-Dimensional Representation of Job History for Economic Adjustment, <a href="#">Federal Committee on Statistical Methodology Conference</a>	2022
Learning Transferrable Representations of Career Trajectories for Economic Prediction, ETH Zurich	2022
Rationales for Sequential Predictions, Google AI NLP Reading Group	2021
Rationales for Sequential Predictions, <a href="#">Hugging Face</a>	2021
<a href="#">Rationales for Sequential Predictions</a> (oral), EMNLP Conference	2021
Text-Based Ideal Points, <a href="#">CFE-CMStatistics Conference</a>	2021
Text-Based Ideal Points, <a href="#">Milstein Program Summer Speaker Series</a> at Cornell Tech	2020
Text-Based Ideal Points, <a href="#">Text as Data Conference</a> at Stanford University	2019
Text-Based Ideal Points, <a href="#">Caselaw Access Project Research Summit</a> at Harvard Law School	2019
<b>Department of Computer Science, Columbia University</b>	
Teaching Assistant, Foundations of Graphical Models (graduate level) Professor: <a href="#">David Blei</a>	2018
<b>Department of Computer Science, Harvard University</b>	
Teaching Fellow, CS 281: Advanced Machine Learning (graduate level) Professor: <a href="#">Finale Doshi-Velez</a>	2015
Teaching Fellow, CS 181: Introduction to Machine Learning Professor: <a href="#">Ryan Adams</a>	2015

Teaching Experience

Conference Reviewing	International Conference on Machine Learning	2017 - 2023
	Neural Information Processing Systems	2017 - 2022
	Advances in Approximate Bayesian Inference	2017 - 2023
	International Conference on Learning Representations	2018 - 2021
	I Can't Believe It's Not Better Workshop	2020 - 2022
	Association for Computational Linguistics	2021
	ACL Rolling Review	2021
	ACM Conference on Fairness, Accountability, and Transparency	2023
Reviewer Recognition	Top 33% Reviewer for ICML	2020
	Top 10% Reviewer for NeurIPS	2020
	ICML Expert Reviewer	2021
	ICLR Reviewer Award	2021
Other Volunteering	<b>Harvard Data Science Review</b> , Early Career Board	2023 -
	<b>Machine Learning in New York City Speaker Series</b> , Organizer	2022 - 2023
	<b>GetUsPPE</b> , Data Scientist	2020
Work Experience	Software engineer intern, <b>Google Brain</b>	2018 - 2019
	Research intern, <b>Facebook Artificial Intelligence Research</b>	2017
	Data science intern (places team), <b>Facebook</b>	2015
	Software engineer intern (data science infrastructure), <b>Facebook</b>	2014
Press	<a href="#">Appearance on Data Skeptic podcast</a> to discuss CAREER	
	<a href="#">Harvard Law Today blog post</a> : Text-based ideal points	
	<a href="#">ProPublica article</a> : Princeton Review price discrimination	
	<a href="#">Today Show segment</a> : Princeton Review price discrimination	