Neil Fasching

Computational Social Scientist | Data Scientist | PhD Candidate
neilfasching@gmail.com | neilfasching.com | 3620 Walnut St. Philadelphia, PA 19104

PROFESSIONAL SUMMARY

PhD candidate at the University of Pennsylvania specializing in computational social science with expertise in experimental design, causal inference, statistical modeling, machine learning, and natural language processing. Highly experienced in designing and executing large-scale studies analyzing social media posts, survey responses, and podcast episodes to understand political behavior and media effects. Research published in top-tier interdisciplinary venues including *Science Advances*, *Journal of Personality*, and *ACL*.

EDUCATION

Ph.D., Communication Expected 2025/2026
University of Pennsylvania

Master's, Statistics and Data Science 2023
The Wharton School, University of Pennsylvania

Master's, Communication Science 2021
University of Amsterdam

SELECTED RESEARCH PROJECTS

Persistent Polarization: The Unexpected Durability of Political Animosity Around Us Elections (Link to Paper)

- Analyzed 66,000 cross-sectional and panel interviews to quantify the effect of election proximity on political animosity
- Employed Interrupted Time Series (ITS) models to causally demonstrate that political animosity remains persistently high and stable around elections, challenging conventional wisdom about polarization

Triangulating Political Toxicity on Twitter (Link to Paper)

- Examined 46.7 million tweets (2012-2022) for toxicity (including harassing, hate, and violent speech)
- Developed scalable pipelines for toxicity measurement using OpenAI and Mistral moderation systems
- Identified patterns in toxicity across time, demographics, and topics using advanced ML models

Partisan Airways: Leveraging Large Language Models to Evaluate Topics of Discussion, Misinformation, and Toxicity on Political Podcasts (Dissertation)

- Analyzed over 28,000 podcast episodes for topics, misinformation, and toxicity
- Developed two novel frameworks for assessing the prevalence of misinformation and hate speech at scale
- Benchmarked LLM-based model performance for transcription, topic segmentation, misinformation identification and hate speech classification

Inconsistent and Very Weak Evidence for a Direct Association Between Childhood Personality and Adult Ideology (Link to Paper)

- Investigated the link between childhood personality and adult ideology using two large longitudinal datasets (combined N=13,822)
- Utilized both frequentist and Bayesian methods (Bayes Factors, q-values), with Bayesian analysis providing strong evidence for the null hypothesis
- Preregistered study on OSF prior to data access to ensure transparent and reproducible research

PROFESSIONAL EXPERIENCE

Computational Research Fellow **Sep 2021 - Present**

University of Pennsylvania

- Collaborate with professors Dr. Yphtach Lelkes and Dr. Duncan J. Watts on projects investigating the influence of news media and social media on human behavior
- Design and execute research studies using diverse data collection, data mining, and analysis techniques
- Analyze large-scale datasets to study trends, patterns, and effects of news media and social media on human behavior
- One example: (mediabiasdetector.seas.upenn.edu/)

Co-Teacher, Modern Data Mining (PhD Level) **Jul 2022 - Present**

The Wharton School, University of Pennsylvania

- Teach PhD-level Data Science course in the Data Science department at Wharton
- Cover cutting-edge machine learning techniques including Boosted Trees, CNNs, RNNs, and LLMs
- Incorporate up-to-date case studies combining statistical theory with practical applications

TECHNICAL SKILLS

Programming Languages: R, Python, SQL, JavaScript

Statistical Methods: Regression (Linear, Logistic, Multilevel), ANOVA, T-Tests, Neural Networks, Ensemble Methods, Time-Series Analysis, Causal Inference, Experimental Design

Research & Analysis: Survey design, experimental design, large-scale behavioral data analysis

Platforms: AWS, Microsoft Azure, Google Colab, Posit Workbench, Git/GitHub

SELECTED PUBLICATIONS

Fasching, N. and Lelkes, Y. (2025). **Model-dependent moderation: Inconsistencies in hate speech detection across LLM-based systems**. In *Findings of the Association for Computational Linguistics*.

Fasching, N., Iyengar, S., Lelkes, Y., and Westwood, S. J. (2024). **Persistent polarization: The unexpected durability of political animosity around US elections**. *Science Advances*, 10(36), eadm9198.

Fasching, N., Arceneaux, K., and Bakker, B. N. (2024). **Inconsistent and very weak evidence for a direct association between childhood personality and adult ideology**. *Journal of Personality*.

Fasching, N., and Lelkes, Y. (2024). **Ancestral Kinship and the Origins of Ideology**. *British Journal of Political Science*.

Pangakis, N., Wolken, S., and Fasching, N. (2023). **Automated annotation with generative AI requires validation**. *arXiv preprint* arXiv:2306.00176.

KEY QUALIFICATIONS

- Extensive experience with survey design, experimental design, and statistical analysis
- Expert in R, Python, and SQL for quantitative analysis
- Experience analyzing large-scale behavioral datasets, including 46.7+ million tweets, 28,000+ podcast episodes, 66,000+ survey interviews, 1.3+ million synthetic sentences
- Published researcher with proven ability to design studies, execute research, and communicate findings
- Causal inference expertise: Interrupted Time Series (ITS), Difference-in-Differences (DiD), Regression Discontinuity Design, Instrumental Variables, Propensity Score Matching