

Philine Widmer · CV

Website: philinew.github.io | E-Mail: philine.widmer@gmail.com

RESEARCH INTERESTS

Public Economics and Political Economy

Traditional, online and social media, algorithms, and their relation to political and market institutions

Econometrics and Machine Learning

Natural language processing, computer vision, and other methods for unstructured data in general

RESEARCH APPOINTMENTS

UniDistance Suisse

Faculty of Business and Economics

Associate Professor · Current position

Paris School of Economics (PSE)

Research Group on Globalization, Political Economy, Trade

Affiliated Researcher · Since 2026

Assistant Professor · 2024–2026

Swiss Federal Institute of Technology (ETH Zurich)

Department of Humanities and Social Sciences

Post-Doctoral Researcher · 2023–2024

EDUCATION

Ph.D. in Economics

University of St.Gallen

Committee: Roland Hodler (supervisor), Ekaterina Zhuravskaya, Reto Foellmi · 2017–2023

With the highest distinction (summa cum laude)

Paris School of Economics, Visiting Ph.D.

Host: Ekaterina Zhuravskaya · 2021–2022

ETH Zurich, Center for Law & Economics, Visiting Ph.D.

Host: Elliott Ash · 2019

Master in Quantitative Economics and Finance

University of St.Gallen

2015–2017 · LGT Capital Partners Award for the best degree in the cohort

Bachelor in Economics

University of St.Gallen

2013–2015

PUBLICATIONS AND INVITED REVISIONS

To access a paper, please click on its hyperlinked title (where available).

The Political Effects of X's Feed Algorithm

With Germain Gauthier, Roland Hodler, and Ekaterina Zhuravskaya · Nature (2026)

Ranked #1 on Altmetric for six consecutive weeks as of April 2026, covered by 100+ news outlets worldwide.

We examine the impact of Twitter's/X's feed algorithm on political attitudes. We randomly assigned U.S.-based users to an algorithmic or chronological feed in an experiment. Users who switch from the chronological to the algorithmic feed shift their political opinions toward more conservative positions. By analyzing users' feeds, we confirm that the algorithm promotes more conservative content. We do not find symmetrical effects among users who switched from the algorithmic feed to the chronological feed.

Media Slant is Contagious

With Clémentine Abed Meraim, Sergio Galletta, and Elliott Ash · Conditionally Accepted at the Economic Journal

We examine the diffusion of media slant. We document the influence of Fox News Channel (FNC) on the partisan slant of local newspapers in the U.S. from 1995–2008. We measure the political slant of newspapers by scaling their articles to Republicans' and Democrats' speeches in Congress. Using channel positioning as an instrument for viewership, we find that higher FNC viewership shifts local newspaper content to the right.

Aligning Large Language Models with Diverse Political Viewpoints

With Dominik Stammach, Eunjung Cho, Caglar Gulcehre, Elliott Ash · Proceedings of the Conference on Empirical Methods in Natural Language Processing (2024)

Large language models such as ChatGPT exhibit political biases. If users query them about political information, they often take a normative stance. To overcome this, we align LLMs with diverse political viewpoints from 100,000 comments written by candidates running for national parliament in Switzerland. Models aligned with these data can generate more accurate political viewpoints from Swiss parties, compared to commercial models.

Text Semantics Capture Political and Economic Narratives

With Elliott Ash and Germain Gauthier · Political Analysis (2024)

Our method is publicly available as a Python Package.

We develop a new text-as-data method. Social scientists' interest in how narratives influence politics or the economy has surged. Yet, it remains challenging to identify narratives in large text corpora. Our unsupervised method achieves this goal by identifying entities and mapping the relations between them.

The Dark Side of the Web: Towards Understanding Various Data Sources in Cyber Threat Intelligence

With Noè Canevascini, Pavel Laskov, Irdin Pekaric, Saskia Laura Schröer · IEEE European Symposium on Security and Privacy Workshops (2025)

We analyze 6.6 million posts, 3.4 million messages, and 120,000 darknet websites to extract cyber threat intelligence. About 20% of collected data contains actionable intelligence, with different platforms serving distinct purposes in the cybercrime ecosystem.

Ministers Engage in Favoritism Too

With Noémie Zurlinden · Journal of Public Economics (2022)

We provide large-scale evidence for favoritism by African cabinet members. Using disaggregated health surveys and novel data on cabinet members' birthplaces (from newspapers and other online text

sources), we compare siblings to get causal estimates. Neonates' and infants' mortality is lower when the current health minister was born in their province.

WORKING PAPERS AND WORK IN PROGRESS

Is Propaganda Front-Page News?

Single-authored

I study the mechanics of subtle censorship in Chinese online news. Using a sample of 1M articles, I construct text-based measures of their alignment with the regime's perspective. More aligned articles tend to feature on their outlet's front-page rather than elsewhere on the website. My analyses point to a censorship strategy differentiating news products between an inattentive majority and an attentive minority.

Who Owns the Online Media?

With Ulrich Matter

Who owns a newspaper is relevant from a political viewpoint, since newspaper ownership comes with benefits beyond the financial returns, such as influence on democratic decision-making. We map detailed ownership networks behind thousands of online news websites in the U.S., Canada, and Europe. We show that ownership matters for content and that outlets tend to have one ultimate owner or, otherwise, highly diversified ownership.

Science under Threat? A Natural Experiment in Economics

With Dominic Rohner and Oliver Vanden Eynde · Under review

We exploit a natural experiment, the U.S. government's "blacklist" of undesirable words released in early 2025, to study whether academics react to political pressure. We find a sharp reduction in the use of banned words among economists at universities relying heavily on NSF funding, particularly for content related to gender, race, and environment.

Towards a Cheap Personalization of Political Messaging with Large Language Models?

With Tiffany Kreuttsch · Draft available on request (double-blind review)

LLMs could enable political microtargeting at low cost. In an experiment, U.S. respondents each viewed one LLM-generated post, randomly personalized to up to 6 demographics. Personalization increases persuasion when matching on party, almost exclusively among Democrats. We show that LLM messages sound systematically Democratic even when targeting Republicans, demonstrating that LLM biases impact personalization effectiveness.

Deep Latent Variable Models for Unstructured Data

With Elliott Ash and Germain Gauthier

Our method is publicly available as a Python Package.

We develop a class of topic models using distributed language representations and document metadata to support statistical inference. The models are estimated in a neural framework with a regularized autoencoder and can be used for exploration, causal inference, and prediction tasks on large, multimodal (e.g., text and images) corpora. For example, we show that polarization since the 1990s in U.S. Congressional speech is driven by framing, not topic choice.

The Neural Ideal Point Model

With Germain Gauthier and Hugo Subtil · Draft coming soon

We develop a deep learning framework to estimate the ideal point models. It is particularly useful when combined with embeddings to estimate ideal points on massive, multimodal, unstructured datasets, such as texts, images, audio, and video footage. We illustrate the model's versatility by analyzing the ideological stance of political advertisement videos on Facebook and Instagram during the 2024 US elections (among other things).

AI Chatbots as a Source of Political Information: Evidence from Benin

With Joel Eke, Eugenia Frezza, Sergei Guriev, and Ekaterina Zhuravskaya · Fieldwork concluded in July 2026

We conduct a randomized controlled trial with university students in Benin around the April 2026 presidential election. Participants are randomly assigned to a control group, a group incentivized to use Meta's Llama, or a group incentivized to use both Llama and a fact-checking LLM. We measure effects on political knowledge, policy attitudes, institutional trust, and voting behavior.

The Politics of Netflix: Political and Social Dimensions of Streaming Content

With Roland Hodler, David Kündig, and Julius Reinhold

We study the political beliefs and gender norms portrayed in streaming content on platforms like Netflix. Combining web-mined data on available titles with computer vision and NLP annotations, we examine whether consumption is polarized by partisanship and gender, and how content supply evolves over time and across geographies.

Gendered Social Media Consumption and its Consequences

With Germain Gauthier, Roland Hodler, and Katharina Kaepfel

We study how content consumption on social media differs across genders and whether these patterns can be changed. In a two-wave field experiment with US-based Instagram users, we combine surveys with individual-level Instagram data exports and randomize an information treatment on how to curate one's own feed. We measure effects on political beliefs and gender norms.

GRANTS

Swiss National Science Foundation

Political Polarization on Streaming Platforms (project partner) · CHF 800K · 2026–2030

Towards a Cheap Personalization of Propaganda with LLMs (as a PI) · CHF 100K · 2024–2025

Social Media, Well-being, and Political Attitudes (PI: Roland Hodler) · CHF 317K · 2023–2027

Mining Narratives from Text (PI: Elliott Ash) · CHF 100K · 2019–2020

London Business School

RCT: AI Chatbots in a Low-Information Environment (Benin) (co-PI) · GBP 20K · 2026

Paris School of Economics

Seed Funding: Gendered Social Media Consumption (co-PI) · EUR 2K · 2025–2026

Basic Research Fund (GFF) at the University of St.Gallen

Stipend fully funding academic stay at the Paris School of Economics · CHF 50K · 2021–2022

Open Philanthropy

Research Project: Political Risks Posed by Advances in Artificial Intelligence (as a PI) · USD 14K · 2021–2022

PRIZES

Royal Economic Society

The Economic Journal Referee Prize 2025

One of 20 referees selected from 2,800+ reports for exceptional contribution in academic peer reviewing · 2025

LGT Capital Partners

Award for the best M.A. degree in Quantitative Economics and Finance in 2017

CHF 5K · 2017 · University of St.Gallen

TEACHING

UniDistance Suisse

Natural Language Processing and Generative AI · Bachelor and Ph.D. · Since 2026

Recherche Empirique et Econométrie · Bachelor · Upcoming spring 2027

Data Management · Bachelor · Upcoming fall 2027

Paris School of Economics

Agentic AI · Since 2026

Political Economy of AI · Since 2025

Media Economics · 2024–2025

ETH Zurich

Computational Social Science with Images and Audio · 2023

Summer Institute in Computational Social Science · 2021

Guest courses

LMU Munich, USI Lugano · *Text-as-Data for Economists* · 2023, 2024, 2026

University of Hohenheim · *The Use of Artificial Intelligence in Economic Research* · 2024

University of St.Gallen

Teaching assistant for various courses · 2017–2022

PROFESSIONAL SERVICES

Refereeing Activities

American Economic Review, Economic Journal, European Economic Review, Journal of Comparative Economics, Journal of Population Economics, Political Analysis

Seminar and Workshop Organization

PEPES Seminar · Co-organizer on PSE side · 2025–2026

MPWZ-CEPR Text-as-Data Workshop · Half-yearly online workshop co-hosted with ETH, Monash, Warwick, and CEPR

EMPLOYMENT

University of St.Gallen

Swiss Institute for International Economics and Applied Economic Research

Ph.D. Student and Teaching Assistant · 2017–2023

Junior Research Assistant · 2015–2017

McKinsey & Company

Social Sector and Healthcare Practice

Intern (with offer for permanent position) · 2015–2016 · Zurich