Cultural Change Through Writing Style: Gendered Pronoun Use in the Economics Profession*

Camilo García-Jimeno[†1] and Sahar Parsa[‡2]

[1]Federal Reserve Bank of Chicago
[2]New York University

March 19, 2025

Abstract

Through their writing, people often reflect their values. Since the 1970s, academic economists have gradually changed their third-person pronoun choices, from using the masculine form to incorporating feminine and plural forms. We document this transition empirically, and examine the role of social interactions among economists in driving the cultural change reflected in these choices. Our analysis relies on a model where writing style depends on the influence of academic peers, the implicit negotiation between co-authors, and individual authors' preferences for expressing gender equality values in their writing. We directly measure peer influence relying on time-varying academic connections between economists, and propose a methodology that uses a homophily-based model of co-authoring decisions to isolate the effect of peer influence from unobserved personal preferences. The model allows us to decompose the observed changes in writing style over the last 50 years into generational shifts, the increasing prevalence of co-authorship in the profession, the increasing share of female economists, and peer influence. Generational changes and the growing share of women in the profession play a minor role. Early on, contrarian economists accelerated the pace of change in writing styles by moving away from their peers' behavior. The large fraction of conformists and the overall homophily in co-authoring, in contrast, slowed the adoption of innovative writing styles by restricting economists' exposure to peers with different gender-attitude signaling preferences.

Keywords: gender, social norms, co-authorships, social networks

JEL Codes: D71, D83, D85, J16, Z1

*We especially thank Doris Pan, whose research assistance has been extraordinary. We also thank Ryan Perry, Liam Puknys, Bea Rivera, Aryan Safi, and Fangzhou Xie for their help at various stages, and Frank DiTraglia, Diego Jimemez Hernandez, Martin Osborne, Debraj Ray, Ariel Rubinstein, Judit Temesvary, and participants at the UPenn micro seminar, the Northwestern Kellogg Political Economy seminar, the University of Calgary 2024 Empirical Microeconomics Workshop, the 2024 Petralia Political Economy Conference, and the 2024 Fed System Econometrics conference for their suggestions. Finally, we thank Professor Yang Feng for sharing his code with us. The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of Chicago or of the Federal Reserve System.

[†]cgarcia@frbchi.org, 230 S. LaSalle Street, Chicago, IL.

[‡]sahar.parsa@nyu.edu, 710 19th West 4th street, New York, NY 10012.

1 Introduction

The second half of the 20th century was transformed by the rise in female labor force participation, and by what Goldin (2006) calls a "quiet revolution" in workplace attitudes and expectations regarding gender. How does such large-scale cultural change occur? In this paper we argue (and document empirically) that a key driver of the diffusion of cultural innovation, and in particular of the adoption of new beliefs and behaviors related to equality between the sexes, is peer influence within professional social networks.

We study cultural innovation in the context of language use, which prominently reflects societal changes in norms and expectations. In particular, we study the choice of gendered pronoun forms in academic writing as a window into the forces shaping beliefs and behaviors related to gender equality (Baron, 1986). We focus on the social network of academic economic theorists between 1970 and 2019. This is for two reasons. First, the field of economic theory uses mathematical models with abstract agents, where gender *per se* is seldom relevant. This allows us to measure writing style choices directly, as authors can freely choose third-person pronoun genders for the agents in their models.[1] Second, we observe rich individual and professional network data, in a professional environment that is cohesive enough for many authors to know each other directly, yet large enough for many to be connected only indirectly. This is crucial for our empirical strategy aimed at identifying peer influences, which will demand the presence of authors who may influence others only indirectly.

The writing style in Economics has undergone a steady transformation in the past 50 years, seeing an increased variety of third-person pronoun forms. We classify all economic theory papers published between 1970 and 2019 into four groups based on their use of third-person pronouns when referring to the agents in their models: those exclusively using masculine (*he, him, himself*), feminine (*she, her, herself*), plural (*they, them, themselves*), or a mix of two or more pronoun forms.[2] Figure 1 shows a rapid decline in the use of masculine pronouns in economic theory papers, from a dominant position – shy of 80 percent of all published papers in 1970 – to just over 40 percent by 1990, and further down to 20 percent by 2019. Alternative pronoun forms gained popularity at different times and rates. The plural and mixed forms began rising in the mid-seventies, while the feminine form emerged around the 1990s. By 2019, papers using masculine, feminine, and mixed pronouns each constituted roughly 20 percent, with the plural form experiencing a resurgence around the

[1] For instance, in principal-agent models authors may use masculine pronouns for the principal, and feminine pronouns for the agent. Stevenson and Zlotnick (2018) highlight a similar consideration in the context of economics textbooks, noting that authors have greater flexibility in arbitrarily assigning genders to fictional characters.

[2] This includes both the use of grammatical forms such as *he/she*, or the use of different forms to refer to different antecedents.

mid 2010s after two decades of stagnation.

In this paper we study this transition to quantify the importance of social interactions within the professional network in driving the aggregate patterns from Figure 1 and the cultural change they reflect. We do so relying on a discrete choice model of pronoun form, to which we add three novel considerations. First, authors' preferences can depend on an idiosyncratic component common across all their publications, and on a time-varying social component capturing the influences mediated through the underlying professional network. These are the main drivers of the authors' aim at signaling or expressing particular values or beliefs through their choices. We measure the social component using the distribution of writing style choices of past co-authors and citees (their peers). Second, we allow for heterogeneous peer effects through a time-invariant, author-specific coefficient on this social component, such that authors may be *conformists* who move towards their peers' past choices, or *contrarians* who move away from their peers' past choices. Third, we capture the potential conflict between co-authors and the implicit bargain that occurs in such cases, by allowing the choice-specific payoffs for co-authored articles to depend on a weighted average of the individual payoff components of the co-authors.

To illustrate the importance of some of these components, consider the discussion between two prominent theorists, Martin Osborne and Ariel Rubinstein, in the preface to their first edition textbook on Game Theory where they discuss their disagreement over how to handle gendered third-person pronouns (Osborne and Rubinstein, 1994). Rubinstein argued in favor of the masculine form, *he*, which he considered neutral, and that other alternatives would be distracting. In his own words, "... in academic material, it is not useful to wave [language] as a flag." Osborne had a different stance, arguing that using the masculine form would not express neutrality, and that the use of the masculine form to refer to individuals of unspecified gender had sexist origins. He went further, advocating the feminine form, *she*, to refer to all individuals, partly to influence the writing practices of future economists.

Separately identifying the role of social influences from changes in the distribution of preferences (entry of new cohorts) and collaboration in driving the change in writing norms poses several empirical challenges. Ignoring co-authorships, for example, would be a major source of confounding because within-author changes in pronoun choices across articles could reflect co-authors' preferences. We, however, aggregate the co-authors' preferences in co-authored papers using utility weights that depend on pairwise characteristics of the authors such as their differences in seniority, citations, or productivity, plausibly related to the authors' relative influence. This is important because around sixty percent of articles in our data are co-authored.

Another empirical challenge is that author-specific preferences for the different pronoun

forms likely depend on the author's social influences. Although we observe more than one publication for many authors, we cannot simply difference out these author-specific effects because our model is non-linear – writing-style choices are discrete. These author-specific effects are nuisance parameters from the point of view of recovering the peer effects. To address this challenge we rely on the well-documented importance of homophily in academic collaborations: when homophily is an important driver of co-authorship decisions, then homophily along ideological or value-based dimensions may also influence collaborations. Thus, controlling for other observable sources of homophily, observed co-authorship decisions contain information about these latent preferences.

To recover a proxy for these latent preferences, we estimate a statistical model of co-authoring decisions borrowed from the Network Science literature on *community detection* (see Karrer and Newman (2010); Newman (2018)). Community detection models exploit the pattern of observed edges to classify nodes into their most likely community when homophily along a latent type drives the existence of edges between nodes in a network. Subsets of nodes with a large number of edges between them are likely to belong to the same underlying community, defined by the characteristic along which homophily in link formation is present. In our setting, it is natural to think of two latent communities (types); a more liberal one with preferences similar to those expressed by Osborne, and a more conservative one with preferences similar to those expressed by Rubinstein. The workhorse inference-based model of community detection is called the stochastic block model (see Feng et al. (2023)). We use it to classify authors into two communities, modifying it to operate on top of an underlying *acquaintance network* capturing co-authorship feasibility inspired by a related idea in Fafchamps et al. (2010).

We construct this acquaintance network borrowing ideas from Natural Language Processing (NLP). The *word2vec* algorithm generates vector representations (embeddings) of words based on the relative frequencies with which words appear near each other in text (Mikolov et al., 2013). These embeddings can then be used to measure semantic similarity. In a close analogy, which we refer to as *author2vec*, we use co-authorships and cross-citations to learn vector representations for each author, allowing us to measure their academic similarity. We then classify pairs of authors as linked in the acquaintance network if they are sufficiently close in this embedded academic space.

Our community detection model classifies 44 percent of authors in one group (which we refer to as more liberal as it includes Martin Osborne), and 56 percent in the other (which we refer to as more conservative as it includes Ariel Rubinstein). Men and women are equally represented in both communities. Besides allowing us to classify all authors into two groups, our estimates of the community detection model reveal strong homohily in co-authorship

between authors who share ethnic background, gender, sub-fields of specialization, similar ages, and similar citation counts. Similar levels of overall productivity as measured by lifetime numbers of published articles, in contrast, do not predict increased collaboration.

Armed with the authors' community assignments we return to the writing style model and use the assignment to directly control for time-invariant author-level preferences. In the presence of time-varying shocks common to social connections in the co-authoring or citations networks, however, time-variation in the writing style of peers may remain dependent with these shocks even after controlling for time-invariant differences in preferences across authors. To address this issue we rely on a control function approach. We estimate the control function using instrumental variables following ideas similar to Jochmans (2023) and Johnsson and Moon (2021). We leverage the exclusion restrictions suggested by the acquaintance network, using changes in the writing style choices of co-authors and citees of the co-authors and citees of a given author, who are themselves not his acquaintances, as a source of variation in his peers' choices. The instruments are strong predictors of the writing style choices of an authors' co-authors and citees.[3]

We find that, overall, 82 percent of authors are conformists: the likelihood of choosing a particular writing style increases with the share of their peers using that style. The remaining 18 percent are contrarians: the likelihood of choosing a particular writing style decreases with the share of their peers using that style. The average peer effects are slightly larger in magnitude for conformists than for contrarians, so the overall average peer effect is positive. Authors classified as more conservative by our community detection model are more likely to be contrarian than the ones classified as liberal: 20 vs. 12 percent. We also find there is considerably more heterogeneity in peer effects among conformists than among contrarians. Our estimates of the distributions of these peer effects are precise, and yield quantitatively similar first and second moments when using alternative parametric forms.

We find men's and women's preferences to be similar. The share of conformists is slightly higher among women than among men, 85 vs. 81 percent, and women choose the masculine form at slightly higher rates than men, particularly when facing female journal editors. Indeed, women were not early adopters of non-masculine pronoun forms; men started the behavioral transformation. This does not mean that the rapid growth in the share of women in the profession was irrelevant; it is possible that early adopters were responding to the

[3]We also highlight why a reduced form IV strategy would fail to recover a well defined treatment effect in a network context with heterogeneous peer effects like the one we study: in the presence of conformists and contrarians, authors whose professional network is composed of mostly conformists are compliers. Authors whose professional network is composed of mostly contrarians, however, are defiers. As we know from the treatment effects literature, IV does not recover a well defined treatment effect in the presence of defiers except in very special cases (Angrist et al., 1996; Dahl et al., 2023). This partly justifies our more structural approach, and is corroborated by our discrete choice model estimates.

growing presence of women in their professional environment.

Older cohorts were both the first to innovate in their writing styles, and the ones that experienced the fastest rate of change. Although younger cohorts do show some decreasing preference for the masculine form, we find surprisingly small cohort differences at their time of entry. The growth in the size of the profession did not change the distribution of preferences much, either. The revolution we observe was one initiated by economists writing in the 1970s and 80s, while later cohorts followed along.

We perform a series of simulations using our estimates to assess the contributions of the different components in driving the observed changes in writing styles. These exercises suggest that co-authoring was not a quantitatively important driver of innovation in pronoun form choices, while peer effects, estimated to be overwhelmingly positive, were a significant drag on the adoption of the more novel writing styles. We also find non-linearities in the long-run distribution of choices as a function of the composition of the population of authors: if conformists were in the minority, an increase in their share would increase the prevalence of feminine and mixed form papers. However, once conformists become a majority, positive feedback makes the initial distribution of choices dominant in the long run and the masculine form dominates. This exercise highlights the sensitivity of the dynamics of cultural change to the underlying distribution of preferences.

We contribute to the literature exploring drivers of long-term cultural change, in particular shifts in attitudes and behaviors towards gender roles. Previous studies have highlighted the impact of technological and economic shocks, political activism, gender imbalances, and even wars in driving such changes (e.g., Akerloff et al. (1996); Alesina et al. (2013); Goldin (2023, 1991); Grosjean and Khattar (2019)). Our study is novel in two ways: we quantify the role of social influences at the individual level within a professional network, and we explicitly model heterogeneity in peer effects, examining the role of contrarians and conformists in the process of cultural diffusion.

Like in other studies (Fryer and Levitt, 2004; Goldin and Shim, 2004; Lieberson and Bell, 1992), we use language as a window into the forces shaping beliefs and behaviors. Although social influences often play a role in changing norms, distinguishing between structural economic considerations and peer influences is challenging. For instance, in the context of women's adoption of their husband's surname, Goldin and Shim (2004) argue that later age-at-marriage, the pill, and increased educational attainment made keeping the maiden name more valuable for women. Studying gendered pronoun choice in scientific publications is convenient, in contrast, because similar economic forces are mostly absent. Indeed, different writing styles do not alter an article's scientific quality, nor to our knowledge currently influence editorial decisions in Economics journals.

Because writing style in relation to gendered pronoun choice is, foremost, driven by signaling or social concerns, our study also contributes to the literature on social norms. Young (2014), for example, argues that "Some norms convey intentions, aspects of personal character, or signal membership in a group. Although the behaviors themselves are of little consequence, they have important reputational implications" (p. 6).[4] In models of fashions, for example, (Bikhchandani et al., 1992; Karni and Schmeidler, 1990; Matsuyama, 1991) there is little or no fundamental value to the action around which a norm emerges. Its value is purely social: conformists value taking the action when more people do, and contrarians value taking the action when fewer others do. In these models, the dynamics of a norm –which may include cycles– depend on the distribution of conformists and contrarians in the population, and on the network structure of their interactions. Our study empirically quantifies these dynamics, examining how the distribution of conformists and contrarians influences the diffusion of writing styles.

Many empirical settings involving social interactions also face the difficulty of distinguishing whether observed conformity to a group is the result of social learning about the inherent value of the action, or of a social concern based on peer pressure. In the case of gendered pronoun choice, this difficulty is not present. And while some social norms are maintained partly through explicit or implicit forms of costly community enforcement such as stigma (Kandori, 1992), in the current professional environment such forms of social enforcement are absent, and thus do not constitute omitted sources of variation that could raise an empirical concern. This of course could change in the future if stigma emerges around the use of some third person pronoun forms.

Our study relates as well to the sociology literature that studies change in opinions and values, and its distinction between cohort effects and period effects as drivers of opinion change (Fernández et al., 2024; Manheim, 1952; Rayder, 1965). Unlike this body of work, we measure actual choices –writing styles– that signal views about gender equality, rather than survey-based opinions. Moreover, our methodology allows us to decompose and quantify the implicit cohort and period effects that drive the attitudinal change we observe, and to micro-found them as being driven by social interactions across members of the profession.

Finally, our paper relates to the broader literature on diffusion of information and behaviors in networks. Well documented examples of these forms of diffusion include the adoption of home computers or hybrid corn (Goolsbee and Klenow, 2002; Griliches, 1957), the spread

[4]Another example close to our study is the literature on naming patterns for Black and White babies in the U.S. born in the late 70s and early 80s, Lieberson and Bell (1992). Fryer and Levitt (2004) find that distinctively Black names in the U.S. arose around the Black power and Civil Rights movements in the middle of the 20th century, and argue that the best model for explaining the timing and spatial patterns in naming is one of identity formation along the lines of Akerloff and Kranton (2000)

of protests and bank panics (García-Jimeno et al., 2022; Kelli and Gráda, 2000), or the spread of drug prescription practices (Iyengar et al., 2011) to name a few. We go beyond documenting the diffusion of a behavior to study the competition between alternative writing styles and the gradual replacement of the standard masculine third person pronoun form for the feminine and plural forms over the span of half a century.

2 The economics profession in the last half century

Beyond writing styles, the economics profession has undergone major related changes in the last half-century. Two significant developments are the increasing share of female academic economists and the rising rates of academic collaboration, possibly related to the information technology revolution.[5] We observe these phenomena within the network of economic theorists we study. Panel (a) in Figure 2 plots the share of papers with at least one female author (in blue), and the share of women publishing (in red). The share of women publishing has experienced a steady increase, from 2 percent of all authors in 1970 to 20 percent by 2020. The share of published papers with at least one female author has increased even faster, from less than half a percent in 1970 to more than 30 percent in 2020.[6] If differences in preferences between men and women are substantial, these sweeping changes in the composition of the body of authors could alone explain a large fraction of the observed shifts in pronoun use from Figure 1. Figure 3 plots trends in the writing-style choices of articles with female authors. Perhaps surprisingly, the writing-style choices of women have followed very similar time trends to those of men.

Trends in co-authoring also reveal momentous change. Overall, 32 percent of articles are single-authored, but the trend has been one of rapid increase in co-authoring. Panel (b) of Figure 2 groups authors into 10-year cohorts based on their first publication, and plots, for each year, the share of co-authored papers by economists from that cohort. While in 1970 less than half of all papers published by the 1970s cohort were co-authored, co-authoring rates among all authors are now close to 90 percent. This trend has affected all cohorts similarly, with the most senior economists showing the largest shift in co-authoring behavior. Whether co-authorship acts as a drag or as a catalyst on the diffusion of new writing styles depends on the nature of these collaborations. Increased collaboration can expose authors to others inclined to use pronoun forms they would not have chosen otherwise. If co-authorship is largely driven by ideology-based homophily, however, the potential for exposure to diverse

[5] As Giuliano and Nunn (2021) point out, variable environments are likely to induce cultural change.

[6] The share of women is lower in theoretical fields compared to applied microeconomics and other areas (Chari and Goldsmith-Pinkham, 2017; Lundberg and Stearns, 2019).

writing preferences will be limited. Figure 4 plots the share of articles using each kind of pronoun form, separately for each 10-year cohort of authors. The trends for only-masculine, mixed, and only-feminine forms move in the same direction for all cohorts. More importantly, the level shifts across cohorts are very small. Indeed, among the earliest adopters of feminine-only and mixed styles we find Duncan Black, Gary Becker, William Baumol, Michael Spence, Thomas Roemer, Sam Bowles, Vincent Crawford and Kenneth Arrow, all members of the earliest cohort in our data. All this suggests both small cohort effects, and that innovation in writing style was led by the older cohort of authors.

The process of cultural change is not seamless. New ideas take time to develop into coherent sets of beliefs, expectations of behavior, or shared norms, and they need to compete with existing ones. Different people may hold different levels of attachment to traditional or innovative beliefs and may have different psychological attitudes towards conforming to or moving away from the beliefs and behaviors of their peers. Identifying social influences in writing styles, however, requires that these sources of persistence at the author-level be limited. Indeed, the panel structure of our data allows us to track individual authors over time. Panel A of Table 2 presents the transition matrix formed by computing the conditional probabilities of switching from using a given pronoun form to every other form, across all author-level sequences of pairs of articles. The diagonal elements are the largest, revealing some persistence. The matrix shows considerable variation in all directions, however, but also some asymmetries. For example, while only 6 percent of authors move from an only-masculine paper to an only-feminine paper, 18 percent move from an only-feminine paper to an only-masculine one. To get a sense of what the behavior reflected in this transition matrix would imply in the long run, in the first row of panel B we report the stationary distribution that would obtain in the limit under the transition matrix from panel A. Overall, the long-run distribution has around a third of only-masculine and of only-plural articles, a fifth of mixed articles, and 12 percent of only-feminine articles. These numbers are not far off what the observed distribution looked like around the mid 2000s.[7] To disentangle the roles of co-authorship, cohort-differences, and social influences driving these aggregate patterns we go beyond these descriptive statistics and estimate a model of writing style choice.

[7]The remaining rows in Panel B present the implied stationary distributions that would obtain from transition matrices restricted to sequences of single-authored papers, or to the papers of authors from each of the 10-year cohorts. We report these transition matrices in Table A.11.

3 The data

In this section we describe the five key components of our data collection. The online appendix contains a more detailed description.

Selection of articles. To construct the sample of economic theory papers, we first identified the set of all papers and authors in Economics and Economics-adjacent fields from 1970 to 2020, using the metadata and full texts from two sources: *Jstor* and *Crossref*.[8] We restricted the initial sample of 710,000 published papers using a multi-step process. After cleaning non-research publications (e.g., Note from the editor, Front Matter), we excluded the articles from 35 journals specializing in econometrics, statistics, or unrelated fields such as operations research. We kept all articles from seven journals specializing in economic theory: the Journal of Economic Theory, the American Economic Journal: Microeconomics, Economic Theory, Games and Economic Behavior, International Journal of Game Theory, Games, and the Journal of Public Economic Theory. For all other published papers, we used their full texts to classify them as (likely) theoretical or not, using a list of microeconomics keywords and a list of econometrics keywords. We then applied a set of inclusion and exclusion criteria based on the frequencies of these keywords.

Next, we excluded articles that lacked a publication date, as well as those published before 1970 or after 2019. We also removed articles for which we could not co-reference any third-person pronouns, and articles without identifiable authors. Finally, we excluded articles with four or more authors, and papers from authors who only ever solo-authored.[9] Our final dataset contains 66,854 articles written by 29,302 unique authors. We assigned unique identifiers to each, building an author-level panel dataset using our final set of articles. Appendix subsection 11.1 contains further details.

Third-person pronouns. With our final sample of papers at hand, we proceed to measure our dependent variable: the gendered pronoun forms used to refer to economic agents in each paper. This requires that we distinguish third person pronoun uses that refer to model agents versus uses for any other reason (e.g., to refer to a real person in the context of a citation). The growing share of women and the growing rates of co-authorship in the profession, in particular, make it important that we do not confound our measures of pronoun use with the increasing occurrence of references to female authors or to groups of collaborators. We tackle

[8] We obtained the *Jstor* data under a data user agreement for the project and the *Crossref* data using the defunct *Crossref* API: https://www.crossref.org/education/retrieve-metadata/rest-api/.

[9] Authors who never co-authored constitute isolated components of the network. Because in the first step of our empirical method we classify authors into two underlying types using information from co-authorship links, there is no information to classify isolated components of the network, and we must exclude them.

this problem using a *co-reference resolution* model, a natural language processing (NLP) tool used to determine the antecedents to which a particular pronoun is referring within a text. It identifies instances where different words or phrases, such as pronouns or noun phrases, point to the same person, place, thing, or concept.[10] We use *Allen NLP*, a state-of-the-art co-referencing neural network model. For each paper, we locate all occurrences of third-person pronouns and extract the surrounding text segment for each occurrence. We then apply the *Allen NLP* model to the segments to obtain the corresponding referenced noun for pronoun.

After mapping each corresponding third-person pronoun to proper nouns in every segment, we keep only instances that refer to a noun in a keyword list of economic agents (see subsection 11.2). This list includes nouns like "individual", "worker", or "agent", to name a few. We made sure to include only gender-neutral proper nouns in this list. Figure A.16 presents the top-50 nouns by frequency of use across all papers. For example, 6.5 percent of third-person pronouns refer to the noun "agent". While *Allen NLP* has an accuracy of at least 75 percent in standard English text, our manual checks suggest an error rate of close to zero at the paper level. After having identified the relevant pronouns, we obtained the counts of masculine, feminine, and plural pronouns in each paper. The distribution of pronoun form counts across articles immediately revealed a striking pattern of mass points at 100 percent masculine, 100 percent feminine, and 100 percent plural, with the remainder of papers, those using a variety of forms, typically showing an even balance of them. This feature led us to classify the articles into four distinct groups: masculine-only, feminine-only, plural-only, and mixed if it used a combination of more than one form.

Co-authoring and citations networks. The metadata for each of the papers in our sample includes information on its authors. Based on these data we built a time-varying co-authoring network dataset encoding as edges the cumulative number of co-authorships between every pair of authors every year between 1970 and 2019. Using *Microsoft's Academic Graph* (MAG), we did a similar exercise to build a time-varying citations network. In contrast to the co-authoring network, the citations network is directed, allowing us to distinguish between backwards (i cites j) and forward (i is cited by j) citations.

Other covariates. We also build a series of covariates. First, we assign sub-fields of specialization to authors based on the *Journal of Economic Literature* (JEL) fields classification. To that end, we select a subset of theory-relevant JEL fields and retrieve generative

[10] For example, in the sentence "The consumer maximizes her utility subject to a budget constraint", a co-reference resolution model can recognize that "her" refers to the noun "consumer".

pre-trained transformer (GPT) embeddings for the terms in each field's description.[11] For each author, we generate sentence GPT embeddings for the titles of his papers and the titles of the papers cited in his papers. We then average these embeddings to compute an author-level embedding, and compute cosine similarity distances between each author and each JEL field. Finally, we assign each author to the three sub-fields closest to him in this embedded space. Second, we classify the ethnic origin of authors using *Namsor*, a commercial software tool that identifies the likely regions of origin of names. Third, we classify the authors' gender using R's *Genderize* package, a probabilistic classifier for first names. Fourth, based on the merged MAG, *Crossref*, and JEL datasets we compute citation counts for each author by aggregating the citation counts across all of his papers.

Acquaintance network. Our full social network comprises 30 thousand economic theorists doing research on a variety of sub-fields over a 50-year period. Naturally, differences in productive years and research areas imply that a typical theorist will only know a small subset of the social network, either personally or through their work. In practice, many pairs of individuals whom we observe neither co-authoring nor citing each other would have never had the opportunity to interact professionally. Identifying these pairs of individuals can provide valuable exclusion restrictions for the purpose of recovering peer effects.

With this purpose in mind, we construct an underlying network of "professional interaction feasibility" that we call the *acquaintance network*. To assign acquaintance edges between pairs of economists that are sufficiently close to each other in "academic" space requires that we can measure academic distance. To do so, we exploit the global patterns of observed interaction across the profession, relying on a methodology based on another NLP tool: *word2vec* (Mikolov et al., 2013). This model is used to analyze semantic relationships between words in a corpus of text. It uses the relative frequencies with which pairs of words appear near each other (right before or after, within a few words of each other, etc.) to assign a high-dimensional vector of real numbers to each word, referred to as the word's *embedding*. The embeddings contain cardinal information about the word's meaning in relation to all other words in the corpus.[12]

Step 1: Embedding authors in academic space. We refer to our methodology as *author2vec*. In close analogy to *word2vec*, the whole set of academic articles stands for the corpus, each article stands for a sentence in the corpus, and the authors and cited economists in each paper stand for the words in a sentence. This allows us to compute the relative

[11] The Appendix reports the list of JEL fields. We retrieved embeddings from the text-embedding-ada-002 model through the *OpenAI* embeddings API. See https://openai.com/blog/new-embedding-models-and-api-updates.

[12] We provide a more detailed description of *word2vec* in Appendix subsection 11.3.

frequencies with which pairs of economists appear near each other across all articles. These relative frequencies then inform the estimation of author-level vector embeddings capturing the relative locations of authors in a high dimensional Euclidean academic space.[13] With the embedding vectors for each author at hand, we then compute a scalar distance measure between every pair of authors. We use the normalized dot product of the embeddings (cosine similarity), which is standard in the literature.[14] Intuitively, consider a pair of authors who has seldom cited each other, or been cited in the same articles, or shared co-authors, and whose co-authors and citees do not overlap across articles either. They will be located far from each other in the resulting embedding space, and we would like to conclude that a co-authorship between them is infeasible.

In panel (a) of Figure 5 we illustrate the variation in academic distance we obtain from our *author2vec* methodology. We do so honing into the local professional network of Ariel Rubinstein and Martin Osborne, who both appear as green nodes. Surrounding each of them, in yellow, are nodes representing their 10 closest economists as measured by our *author2vec* similarity metric. The larger-sized nodes among these represent co-authors. The nodes in tan, in turn, represent co-authors outside of their closest 10. The length of the edges in this figure is proportional to the distance in academic space between economists, and dashed edges represent citation relationships. In addition, we labeled a select few nodes with their cosine similarity with either Osborne or Rubinstein.

Similarity predicts co-authorships and citations: one of the ten closest economists to Osborne is his co-author, and three out of the ten closest to Rubinstein are his co-authors. A large fraction of the ten closest to each are also cited by them. While co-authors of each other, Osborne and Rubinstein have a cosine similarity of only 0.12. This is not surprising: they mostly work on distinct lines of research, and there is no overlap between the sets of economists closest to them. Finally, the figure also reveals that while having fewer co-authors than Rubinstein, Osborne is on average closer to his local network than Rubinstein is. Thus, Rubinstein's professional relationships appear to be more academically diverse. The illustration reassures us that the measure of similarity we computed captures meaningful variation in academic proximity necessary to construct the acquaintance network.

Step 2: Building the Acquaintance Network. Using the academic distances from our *author2vec* methodology, we build an 'acquaintance set' $Q(i)$ for each author i. It is the set of authors with whom, we believe, i could potentially form co-authorships. Effectively, it defines an underlying network on top of which actual co-authorships may form.

[13] We set to 100 the dimension of the author embedding vectors. In *word2vev* and all other GPT models, the dimensionality of the embedding space is a model parameter. Naturally, larger corpora allow for higher dimension embeddings.

[14] Cosine similarities have support in $[-1, 1]$.

We construct $Q(i)$ as follows. We build a neighborhood $L_n(j)$ for each author j, that includes the n closest authors to j based on our cosine similarity measure. We call $Y(i)$ the set of years when author i is active, which we define as the range starting 3 years before his first publication, and 5 years after his last publication. Finally, calling $C(i)$ the set of co-authors of i, including himself, we define an acquaintance set as

$$Q_n(i) = \{k : k \in L_n(j), j \in C(i), Y(k) \cap Y(i) \neq \varnothing\}.$$

The collection of these sets defines an acquaintance network. Our benchmark estimates below use $n = 10$, with alternative specifications using $n = 5$ or $n = 20$. The acquaintance sets aim to include all authors who have co-authored with author i and those who are sufficiently close in academic space to be considered potential co-authors, as long as they overlap in their active years. Indeed, average academic cosine similarity among co-authors, acquaintances, and non-acquaintances are 0.53, 0.41, and 0.03. As an illustration, consider panel (b) of Figure 5, where we plot conditional densities for the cosine similarities between Ariel Rubinstein and all other economists in the professional network. The distribution of similarities with his non-acquaintances (in purple) is centered around 0, considerably to the left of the distributions with his acquaintances (in blue, red, and green), centered around 0.25. While there is a thin right tail of non-acquaintances with almost no density above 0.3, much of the density of acquaintances lies above 0.3. Rubinstein's co-authors are heavily concentrated above the 75th percentile of the acquaintance distributions, and above the 99th percentile of the non-acquaintance distribution. Indeed, three of the ten most similar authors to Rubinstein are also his co-authors (Eliaz Kfir, Michael Richter, and Yuval Salant). These patterns are typical across most economists.

Table 1 also illustrates a pattern of increased homophily on observables as we compare the overall sample of pairs of authors to those with an acquaintance relation, and to those who co-authored. In addition, Figure A.13 presents the cross-scatter plots of co-authorship, forward and backward citations, and acquaintance log degree distributions. Naturally, degree in any one network is strongly predictive of degree in any other. There is, however, wide variation in the acquaintance degree at any level of co-authorship or citation counts, particularly among authors with few co-authors or with a low citation count.

4 Model of writing style

We propose a discrete-choice model of pronoun writing style. The author(s) of each paper decide among four possible styles: masculine-only, feminine-only, plural-only, or a mix of

third person pronouns. We denote a generic choice by $\rho \in \{m, f, p, x\}$. It can be either a joint choice in co-authored papers or an individual choice in single-authored papers. We consider only two-person co-authorships for simplicity, but the model can readily accommodate three or more authors. An author's utility from a choice depends on an author-level, time-invariant component, which we think of as capturing latent ideology or values, and on a social interaction component capturing the influences from his professional network. The nodes in this network change over time as new economists join the profession, and the edges change over time as new publications appear, implying new co-authorships and citations. We will allow for heterogeneity across authors in their response to the social influences, and will impose some structure on the nature of this heterogeneity. In a co-authored paper, the utility of a choice is a weighted average of the preferences of its authors.

4.1 Pronoun-choice payoffs

We denote by $a(ij)t$ an article written by authors i and j published in year t. Without loss of generality, we index single-authored papers as $a(ii)t$. The payoff from choice ρ is

$$u_{a(ij)t}(\rho) = \alpha_\rho + \phi(\mathbf{z}_{ij})[\beta_i r^\rho_{it} + \delta^\rho_i] + (1 - \phi(\mathbf{z}_{ij}))[\beta_j r^\rho_{jt} + \delta^\rho_j] + \epsilon^\rho_{a(ij)t}, \tag{1}$$

where $\phi(\mathbf{z}_{ij}) \in [0, 1]$ and $\phi(\mathbf{0}) = 1/2$ is a bargaining weight. It represents author i's weight in the joint decision, and $\mathbf{z}_{ij}$ is a vector of pairwise covariates such as their difference in ages, difference in citation counts, etc. The bargaining weight is the same across choices, as there is no reason why the relative influence of the authors should differ across the choices. The r^ρ_{it} captures the social influences from i's professional network. We measure it as the fraction of all papers written by his previous co-authors and his previous citees that have used writing style ρ. Formally,

$$r^\rho_{it} = \frac{\sum_{k \in C_i(t) \cup H_i(t)} \sum_{\{a(k\cdot)\tau : \tau < t\}} \mathbf{1}\{\text{Paper } a(k\cdot)\tau \text{ uses } \rho\}}{\sum_{k \in C_i(t) \cup H_i(t)} \sum_{\{a(k\cdot)\tau : \tau < t\}} 1}, \tag{2}$$

where $C_i(t)$ and $H_i(t)$ represent the set of co-authors and citees of author i up to time t. r^ρ_{it} varies across i's publications as his professional network evolves, and across choices as a function of the previous choices of his network. It also varies across economists as each faces a different network of peers. The β_i represent peer effects. These can vary across authors but are constant for an author across his publications. We refer to authors for whom $\beta_i > 0$ as conformists, and refer to authors for whom $\beta_i < 0$ as contrarian. While conformists tend to prefer choices popular among their network, contrarians tend to dislike choices popular

among their network *regardless of the choice*.

To allow for this form of 'psychological' heterogeneity to vary with some author-level characteristics $\mathbf{w}_i$, we introduce two unobserved types $\psi_i \in \{\underline{\psi}, \overline{\psi}\}$ representing conformist and contrarian preferences, and allow the distribution of ψ_i to depend on $\mathbf{w}_i$. Conditional on $\underline{\psi}$ an author draws his β_i from a distribution with positive support. Conditional on $\overline{\psi}$ an author draws his β_i from a distribution with negative support. In practice we draw these coefficients from beta distributions and include a scale factor S_ψ:

$$\begin{aligned}\beta_i|\underline{\psi} &\sim Beta(a_{\underline{\psi}}, b_{\underline{\psi}}, S_{\underline{\psi}}) \\ -\beta_i|\overline{\psi} &\sim Beta(a_{\overline{\psi}}, b_{\overline{\psi}}, S_{\overline{\psi}}).\end{aligned}$$

This captures the idea that a conformist follows his peers regardless of what they are choosing, and a contrarian moves away from his peers regardless of what they are choosing.

We also allow for time-invariant author-specific differences in their preference for the different writing styles, δ_i^ρ. These represent differences in values or beliefs related to the expression of gender equality in writing and are from our point of view, unobserved. The δ_i^ρ are incidental parameters. Despite the panel nature of our data, in a non-linear context like this one we cannot simply difference them out.[15] Based on our motivating discussion on the Osborne-Rubinstein debate, without loss of generality we write these fixed effects as follows: some authors are ideologically similar to Rubinstein, drawing their δ_i^ρ from a distribution with mean δ_R^ρ. Some authors are ideologically similar to Osborne, drawing their δ_i^ρ from a distribution with mean δ_O^ρ. If we define O_i to be a dummy variable indicating beliefs similar to Osborne's, we can write

$$\delta_i^\rho = \delta_R^\rho(1 - O_i) + \delta_O^\rho O_i + \mu_i^\rho, \qquad \mathbb{E}[\mu_i^\rho] = 0 \tag{3}$$

In this setting it is natural to expect the social interaction component to be dependent with these author-specific effects: $\mathbb{E}[r_{it}^\rho \delta_i^\rho] \neq 0$. If there is preferences-based homophily in peer choice, for example, previous choices of peers will be dependent with own ideology. Even in the absence of such homophily, in the presence of peer effects, i's past choices, which depend on his values, may have influenced his peers past choices.

[15] A common approach in discrete choice settings following Chamberlain (1980) is to write down the conditional likelihood of the data, conditioning on a sufficient statistic for the incidental parameters. This sufficient statistic turns out to be the vector of total counts of realized choices across all observations for a given unit. In our setting such an approach has two disadvantages. First, two thirds of the articles are co-authored, so the choice-specific payoffs depend on two different nuisance parameters. Second, a conditional likelihood approach does not allow recovering the contribution of the author-specific effects to the distribution of observed choices, which would restrict our ability to decompose the evolution of writing style norms we observe into the contributions of peer influence, underlying values, and co-authorship.

The α_ρ are choice specific intercepts. Key to our setting, throughout we will maintain that $\mathbb{E}[r_{it}^\rho \alpha_\rho] = 0$. This is, in contrast to much of the empirical literature estimating discrete choice models, ours is a setting without unobserved choice-specific fixed effects that may be dependent with the endogenous regressor of interest. For example, in models of residential location choice (Bayer and Timmins, 2007) or of differentiated product demand (Nevo, 2003), choice-specific payoffs depend on choice-specific unobserved attributes that make a choice more or less valuable to everybody (e.g., access to public transportation in the case of neighborhoods, sweetness in the case of cereals, etc.). The gender choice of third-person pronouns is one where unobserved attributes are absent because the value of a given choice is purely social, this is, it depends only on how much others value it and the signaling concerns around this. The scientific contribution of an economic theory paper is invariant to the choice of gender for the pronouns used in it.[16]

Finally, the $\epsilon_{a(ij)t}^\rho = \varphi_t^\rho + \widetilde{\epsilon}_{a(ij)t}^\rho$ represent time-varying unobservables. On one hand, φ_t^ρ may represent overall trends in relative popularity of writing style ρ coming from outside the economic theorist professional network. We will account for them with time fixed effects. On the other hand, $\widetilde{\epsilon}_{a(ij)t}^\rho$ may represent idiosyncratic shocks affecting the authors of paper $a(ij)t$, possibly dependent with $(r_{it}^\rho, r_{jt}^\rho)$: $\mathbb{E}[(r_{it}^\rho, r_{jt}^\rho)\widetilde{\epsilon}_{a(ij)t}^\rho] \neq \mathbf{0}$. For example, auto-correlation in the $\widetilde{\epsilon}_{a(ij)t}^\rho$'s will generate dependence with the social influences, r_{it}^ρ, through network effects: i's past shock induces him to choose a particular writing style; his conformist peers will subsequently mimic his choice; their choices now influence i at time t.

Even if in the population as a whole the distribution of conformism and contrarianism is stationary over time, the composition of the pool of academic economic theorists may have changed ideologically, and definitely has seen an increasing share of women. Thus, we allow $\mathbf{w}_i$, the characteristics governing the distribution of 'psychological' types ψ_i, to include O_i and sex[17], as it may be that men and women or conservatives and liberals differ in their

[16] An argument could be made that the use of some pronoun forms can improve the readability or the quality of the writing in a paper, implying differences in the fundamental value of the different choices. For example, giving different genders to different players in a model –the mixed choice–, may allow the reader to more easily identify who the author is referring to. The recent rise in popularity of the feminine-only and the plural-only forms suggest many authors do not share this view. Moreover, some may argue that this added flexibility could easily make writers less careful in constructing their sentences, and may thus hurt the writing quality itself. Thus, we believe that any differences in scientific writing quality directly coming from using one type of pronoun form over another are at most second order. Another possibility is that differences in the fundamental value of using different pronoun forms emerge from editorial practices that condition the likelihood of publication of a paper on the pronoun form used in it. We are unaware of such behaviors in the Economics profession. Even if editors do no condition their publication decisions on pronoun choice, this could still be a concern if an author holds the (wrong) equilibrium belief that they do. In subsubsection 6.2.1 below we empirically test this possibility.

[17] Because we use gender to refer to the grammatical category for third person pronoun choices, to avoid confusing the reader here we use the word sex to indicate whether the authors are men or women.

psychological inclination toward conformism or contrarianism. Finally, only-plural will be the baseline category for estimation.

To estimate the discrete choice model based on the preferences in (1) we must address two main econometric challenges. First, idiosyncratic time varying unobservables, $\widetilde{\epsilon}^{\rho}_{a(ij)t}$, may be dependent with social influences through the network structure. Second, ideological preferences, O_i, are unobserved and may also be dependent with social influences. We now address the first concern, and turn to the second concern in section 5.

4.2 Identification: leveraging the acquaintance network

The acquaintance network we described in section 3 will provide us with the exclusion restrictions to address the empirical challenges highlighted above.

4.2.1 Control function

We first address the dependence between social influences and time-varying unobservables. We do so with a control function approach following ideas in Jochmans (2023) and Johnsson and Moon (2021), and exploiting our acquaintance network. Specifically, we build time-varying instruments for r^{ρ}_{it}. The logic of these instruments is as follow: if an author i has peers (past co-authors and citees) who are themselves subject to peer influences, then the pronoun choices of the peers of these peers will generate variation in their pronoun choices, this is, in r^{ρ}_{it}. The choices of past peers of an authors' peers are thus relevant. If we can find peers of i's peers who are *not* in i's acquaintance set, then we know they do not directly influence his writing style. The choices of non-acquaintance peers of i's peers are thus excludable. We construct such instruments as

$$z^{\rho}_{it} = \frac{\sum_{k \in C_i(t) \cup H_i(t)} \sum_{P_i(k,t)} \mathbf{1}\{\text{Paper } a(\ell m)\tau \text{ uses } \rho\}}{\sum_{k \in C_i(t) \cup H_i(t)} \sum_{P_i(k,t)} 1}, \tag{4}$$

where $P_i(k,t) = \{a(\ell m)\tau : \tau < t \text{ and } \ell \in C_k(t) \cap Q^C_i, m \in Q^C_i\}$ denotes the set of articles by authors who are not acquaintances of author i, but who are past co-authors of one of his past co-authors or citees, k.[18] Because ours is a panel data setting, we use variation over time in the indirect exposure to non-acquaintances of peers' writing-style choices. With this

[18] If the set $\bigcup_{k \in C_i(t) \cup H_i(t)} P_i(k,t)$ is empty, this is, if for a given author-publication period none of his co-authors or citees have co-authors that are not his acquaintances, we define $z^{\rho}_{it} = 1/4$ for all $= \rho$, the maximum entropy multinomial distribution among four choices.

purpose in mind, we compute first difference versions of (4), $\Delta z_{it}^{\rho} \equiv z_{it}^{\rho} - z_{it-1}^{\rho}$, where the difference is between two consecutive papers published by author i.

To understand what the variation in Δz_{it}^{ρ} is capturing, consider Figure 6 as an illustration. In it, we represent the professional network of Debraj Ray across two years when he published papers: 1993 on the left and 1994 on the right. Green circles represent his acquaintances, and the name labels represent people he had co-authored with prior to the corresponding year. Pink circles represent the co-authors of Ray's co-authors who are not in his acquaintance set. For example, in 1993 Douglas Bernheim had 4 co-authors who were themselves outside Ray's acquaintance set. Between 1993 and 1994, Ray developed several new co-authorship relationships, one of them with Kalyan Chatterjee. Chatterjee had himself three past co-authors who were not Ray's acquaintances. Thus, the past writing style choices of these three people generate the variation in the instrument for Ray in 1994. Figure 7 provides a similar example, this time looking at Drew Fudenberg's professional network. Between 1992 and 1993 he developed a new co-authorship relation with Christopher Harris, and Harris had himself three past co-authors not in Fudenberg's acquaintance set.[19]

If the distances in academic space we measured relying on our *author2vec* methodology do capture the relative professional proximity of authors –e.g., feasibility of co-authoring, visibility for citing–, these instruments will be valid.[20] Because the social interaction variables r_{it}^{ρ} in (1) are fractions that add up to 1 across ρ, we implement our control function approach by estimating a fractional response multinomial logit reduced form regression (Mullahy, 2011), under which:

$$\mathbb{E}[r_{it}^{\rho}|\Delta \mathbf{z}_{it}] = \frac{\exp(\Delta \mathbf{z}_{it}'\boldsymbol{\pi}^{\rho})}{1 + \sum_{\rho \in \{m,f,x\}} \exp(\Delta \mathbf{z}_{it}'\boldsymbol{\pi}^{\rho})}, \tag{5}$$

where $\Delta \mathbf{z}_{it} = (\Delta z_{it}^{m}, \Delta z_{it}^{f}, \Delta z_{it}^{x})$. These conditional mean functions capture the part of the variation in i's peer writing-style choices induced by time-series variation in the choices of their peers who are themselves non directly connected to i. Under the identifying assumption that $\mathbb{E}[\Delta z_{it}^{\rho} \widetilde{\epsilon}_{a(ij)t}] = 0$, the residuals from the fractional response model, $\eta_{it}^{\rho} = r_{it}^{\rho} - \widehat{\mathbb{E}}[r_{it}^{\rho}|\Delta \mathbf{z}_{it}]$, contain the endogenous variation in r_{it}^{ρ} which we include as a regressor in (1). Notice that

[19] Jochmans (2023) proposes a similar approach, in the context of endogenous selection of peers. There, link decisions that involve a given individual do not need to be independent of one another, but they are independent of link decisions made by other pairs of individuals located sufficiently far away in the network. This is different from the popular IV strategy from Bramoulle et al. (2009) which uses covariates from second-degree neighbors as instruments in a cross-sectional setting, and does not require the absence of alternative paths between those second-degree neighbors and a given node. Because we have access to panel data, we can use past choices instead.

[20] Notice that the exclusion restrictions provided by the acquaintance network imply that $\mathbb{E}[z_{it}^{\rho} \mu_i^{\rho} | O_i] = 0$.

in co-authored papers, we need to include both η_{it}^{ρ} and η_{jt}^{ρ}[21].

Table 3 reports our estimates of the the $\boldsymbol{\pi}^{\rho}$ coefficients from (5). Throughout, the baseline category is only-plural. The top panel reports our benchmark results, where we consider both co-authorships and citations as generating edges in the professional network. The columns report, in that order, results for the share of only-masculine, only-feminine, and mixed papers among an author's peers. Reassuringly, and relative to the only-plural share, the largest coefficient in each equation is the one for the instrument of the corresponding pronoun form. These first stages yield highly statistically significant coefficients. The middle and bottom panels explore the robustness of these results to restricting the professional network to only co-authorships, or only citations. In both cases the pattern of results is unchanged, with somewhat larger coefficients when restricting attention to the citations network.

Table 4 explores further the robustness of these result, this time in relation to functional form. Rather than a fractional multinomial logit model, we present results from linear regression models separately for each pronoun form share. Instead of using the first difference of the instruments as regressors, we use their levels from (4), and include author-level fixed effects to exploit only within-author variation. The results are consistent with those from Table 3: the instrument for the corresponding pronoun form always positively and significantly predicts the peer's pronoun form share. Just as before, we present results for networks that include both co-authors and citees in the top panel, only co-authors in the middle panel, and only citees in the bottom panel. These results also mimic those from Table 3.

4.2.2 Discussion: what about linear IV?

The availability of instruments for the endogenous pronoun form choices of peers we just described may raise the following question: why not estimate a linear two-way author-fixed effects model where the dependent variable is a dummy for a given pronoun form choice, and we instrument the peers' average choices with the same instruments we use to estimate the control function outlined above? Such an IV estimator would not in general identify any well defined causal effect in a setting like ours. This is because in a network setting with heterogeneous peer effects, the standard monotonicity requirement for the first stage will not hold. In the treatment effects literature, for example, it is well know that IV in general does not recover a treatment effect for any sub-population in the presence of defiers (Angrist et al., 1996; Dahl et al., 2023).

In our setting, we have argued there can be conformist and contrarian economists. Consider an author i, conformist or contrarian, with one peer, j. If j is a conformist, he will be

[21] Johnsson and Moon (2021) also propose using a control function to recover peer effects, but do so in a cross-sectional network setting with endogenous network links instead.

more likely to choose a particular writing style when that style is more popular among his peers, k. Author i is thus a *complier*, since his treatment variable, j's choice, is increasing in k's choice. If j is a contrarian, however, he will be more likely to choose a particular writing style when that style is less popular among his peers, k. Author i is thus a *defier*, since his treatment variable, j's choice, is decreasing in k's choice.

5 Co-authorship formation model

In this section we turn to addressing the second econometric challenge we raised in subsection 4.1, namely the presence of unobserved time-invariant preferences related to author's values or beliefs, possibly dependent with the social influences they experience. Our starting point is the ample empirical evidence of homophily in academic collaboration networks. Numerous studies have explored this question in the context of the Economics profession. Besancenot et al. (2017); Ductor et al. (2023); Ductor and Prummer (2023); Freeman and Huang (2014); Önder et al. (2021), for example, show there is strong homophily in co-authorship along the gender, ethnicity, fields of specialization, and productivity dimensions. Fafchamps et al. (2010) argue, in addition, that social distance between economists also drives co-authorships. In work studying other scientific fields, Combes and Givord (2018) also find strong gender homophily in co-authorships, and Newman (2001) finds that network distance (what he refers to as 'small worlds') is also a driver of collaborations.[22]

If there is homophily in co-authorships along the unobserved preferences reflected in δ_i^ρ, we would expect clustering of edges: a high prevalence of co-authorship between pairs of authors with relatively more conservative (or traditional) preferences ($O_i = 0$) and between pairs of authors with relatively more liberal (or innovative) preferences ($O_i = 1$), compared to pairs of authors with dissimilar ones. Our key observation is that observed patterns of co-authorship across the whole network of economic theorists will contain information about ideological similarity. Consider, for example, a setting where there is homophily in collaborations based on an observed and an unobserved dimension. Suppose two authors in this network are very dissimilar in the observed characteristic but are, nevertheless, seen co-authoring. We may infer they are similar in their unobserved characteristic. Conversely, suppose two authors

[22]Recent work has also explored potential reasons for the observed homophily along gender lines, including preferences over risk (Lindenlaub and Prummer, 2020), gender-based asymmetries in recognition for collaborative work (Sarsons et al., 2021), and signaling concerns (Onucich and Ray, 2021). Related work has shown that the extent of co-authorship has increased substantially in the last half a century (Anderson and Richards-Shubik, 2022; Hammermesh, 2013; Kuld and O'Hagan, 2018; McDowell and Melvin, 1983), that the average age of authors has been increasing (Hammermesh, 2015), and that the Economics profession shows "small world" patterns because a small number of star authors co-author widely with other authors who themselves have few co-authors (Goyal et al., 2006).

are very similar in their observed characteristic but are not seen coauthoring. We may infer they are dissimilar in their unobserved dimension.

5.1 The community detection problem

Based on the observation above, we estimate a simple homophily-based model of co-authorship borrowing ideas from the Network Science literature.[23] This literature has studied extensively the problem of detecting unobserved types within a network. It refers to it as the *community detection* problem (Karrer and Newman, 2010; Newman, 2001, 2018). This will allow us to classify all authors in the network into two groups, or communities. One which we interpret as having preferences more similar to those expressed by Rubinstein, and one with preferences more similar to those expressed by Osborne. Thus, we will be able to recover an estimate of O_i for all authors.[24]

The workhorse inference-based model for community detection is called the Stochastic Block Model (SBM). Here we implement a covariates-adjusted SBM following Feng et al. (2023), to allow for homophily in other observable characteristics. The SBM presumes the existence of a finite number of communities, which we will fix to two, takes a set of nodes –authors in our setting–, and models the number of links between every pair of them as draws from a Poisson distribution that depends on pairwise characteristics. Such a model is ideally suited to settings where the network is sparse, and where there can be more than one edge between a pair of nodes. Both of these are features of co-authoring network: the number of co-authorships is very small compared to the number of all potential co-authorships, and some pairs of authors share more than one publication.[25]

We modify the workhorse SBM model in one key way: whereas the standard model allows for edges to form between every pair of nodes, we restrict co-authorships to arise only among pairs of authors who are in each others' acquaintance sets. This is for two reasons. First, an econometric challenge that arises in models estimating dyad-level equations is the quadratic explosion of potential edges as the the number of nodes grows, introducing computational

[23] In Economics, several papers have proposed approaches to estimate models of network formation with unobserved drivers of link formation. In some instances these models bypass estimating the unobserved effects (Fafchamps et al., 2010; Graham, 2017). In others, they exploit additional structure (dePaula et al., 2018; Islam et al., 2022).

[24] The SBM relies on the global patterns of co-authorship to infer the community memberships. Because authors who never co-authored are isolated individual components of the overall graph, there is no information to classify them into either community. Thus, we must exclude them from our estimating sample. 11 percent of authors in our original network never co-authored, of whom 87 percent only published one paper.

[25] Another convenient feature of the SBM in our setting is that it allows us to abstract away from the time dimension of the co-authoring problem: it models the intensive margin of co-authorship. Because our interest is to learn a time-invariant feature of the set of authors, the approach is not very restrictive while allowing us to avoid modelling the complex dynamics of co-authoring decisions over time.

difficulties. Second and more importantly, in relatively large social networks a large (and increasing in network size) share of potential edges are infeasible because most individuals, in practice, can only form relationships with a local subset of others. Considering dyads who could not possibly have formed a connection will lead to bias in the estimates of the strength of homophily. Suppose, for example, that shared ethnicity does increase the likelihood of link-formation, but that only geographically close people are feasible candidates for collaboration. If pairs of people who would never face the opportunity of collaborating share, for example, the same ethnicity, a model that uses the information from these pairs of people will underestimate, possibly severely, the importance of ethnic similarity. Because we constructed the acquaintance network precisely to capture subsets of authors who are likely to be near each other in the professional sphere, restricting the possible co-authorships to happen only between mutual acquaintances allows us to address this issue. The acquaintance network is, thus, an underlying set of edges on top of which co-authorships can be formed.

5.2 The acquaintance network-adjusted SBM

Each of the n economic theorists has an unobserved (to us) type $\tau_i \in \{\ell, c\}$. These two types differ in their ideological preferences. The fraction of ℓ types in the population is π_ℓ, and the fraction of c types is $\pi_c = 1 - \pi_\ell$. Conditional on types, the number of co-authorships between $i \in Q(j)$ and $j \in Q(i)$ is Poisson distributed:

$$a_{ij} \sim \mathcal{P}(\omega_{\tau_i \tau_j} e^{\mathbf{x}'_{ij}\boldsymbol{\gamma}}),$$

where $\boldsymbol{\tau} = (\tau_1, \tau_2, ..., \tau_n)$ is the vector of true types, and

$$\Omega = \begin{pmatrix} \omega_{\ell\ell} & \omega_{\ell c} \\ \omega_{\ell c} & \omega_{cc} \end{pmatrix}$$

governs the degree of type-based homophily in the network formation technology.[26] Diagonal entries that are larger in magnitude than the off-diagonal entries reflect homophily. We do not impose such a constraint in estimation. To allow for homophily along observables, we include the following pairwise covariates in $\mathbf{x}_{ij}$: dummies for same ethnicity and same gender, a count variable for the number of common sub-fields, the author's age difference, the author's log citations difference, the author's log productivity difference, and the log of the product of the author's productivities[27].

[26] The model accommodates single-authored papers in the form of 'self-edges'.

[27] Newman (2018) shows that including this last covariate is akin to a SBM with 'degree-correction', this is, it accommodates networks with high dispersion of its degree distribution.

The joint likelihood of observing co-authoring matrix $\mathbf{A}$ and an assignment of types $\boldsymbol{\tau}$ is

$$\begin{aligned}\mathcal{L}(\mathbf{A}, \boldsymbol{\tau}|\Omega, \boldsymbol{\gamma}, \pi, \mathbf{X}) &= \mathbb{P}(\mathbf{A}|\boldsymbol{\tau}, \Omega, \boldsymbol{\gamma}, \pi, \mathbf{X})\mathbb{P}(\boldsymbol{\tau}|\Omega, \boldsymbol{\gamma}, \pi, \mathbf{X}) \\ &\propto \prod_{i=1}^{n} \left[\prod_{j \in Q(i)} \left(\omega_{\tau_i \tau_j} e^{\mathbf{x}'_{ij}\boldsymbol{\gamma}} \right)^{a_{ij}} \exp(-\omega_{\tau_i \tau_j} e^{\mathbf{x}'_{ij}\boldsymbol{\gamma}}) \right] \pi_{\tau_i}. \end{aligned} \tag{6}$$

Solving the community detection problem entails maximizing this likelihood jointly over the type shares, π_ℓ, the homophily coefficients $\boldsymbol{\gamma}$ and Ω, and the assignment vector $\boldsymbol{\tau}$. In Appendix subsection 11.6 we provide the details of the procedure, which closely follows Feng et al. (2023). It entails noticing that the maximum likelihood estimates for π_ℓ and Ω have closed forms as functions of $\boldsymbol{\tau}$ and $\boldsymbol{\gamma}$ only, so we can compute a profile likelihood that is only a function of $\boldsymbol{\tau}$, and $\boldsymbol{\gamma}$. In turn this profile likelihood can be optimized in two steps. First, finding the optimum over $\boldsymbol{\gamma}$ for a fixed $\boldsymbol{\tau}$. Then, optimizing over $\boldsymbol{\tau}$ using an EM algorithm.

5.3 Estimation results form the co-authoring model

Table 5 presents our estimates. The first column reports our benchmark results, under the acquaintance set definition based on the ten closest economists as we described in section 3. The second and third columns report results under tighter (five closest) and looser (twenty closest) alternative acquaintance set definitions. The top block reports estimates for $\boldsymbol{\gamma}$. Except for the difference in lifetime productivities between the pair of authors, all other pairwise covariates are strong predictors of co-authorship. Pairs of economists of the same ethnic background, same sex, and sharing sub-fields in common are more likely to write together. Larger age gaps and larger citations gaps decrease the likelihood of co-authoring. The point estimates on all these pairwise covariates are very precisely estimated, confirming strong homophily in academic collaboration within the economic theory field.

The bottom block reports the estimates for implied homophily along the unobserved type dimension, Ω. These parameters are informed by the relative frequencies of observed co-authorships given the optimal community assignment $\boldsymbol{\tau}$. They reveal, for example, that conditional on observables, average co-authorships are nine times higher between ℓ types than between an ℓ and a c type, and three times higher between c types than between an ℓ and a c type. These parameters are also very precisely estimated.

Turning to the community assignment $\boldsymbol{\tau}$, we classify 56 percent of authors into one group, and 44 percent into the other. In the absence of external information the assignment cannot tell us what each group represents, which is the standard label switching issue. In our context however, we do have additional information. Across all specifications, we always find Martin

Osborne and Ariel Rubinstein to be classified in different communities. As a result, we refer to the two groups as the (relatively more innovative) Osborne type, and the (relatively more traditionalist) Rubinstein type. The Rubinstein group is the largest of the two.[28] Indeed, across the two alternative specifications in the table, 85 and 82 percent of authors are consistently classified into the same group as they are in the benchmark specification.

While overall we find relatively balanced group sizes, the growth in the profession since the 70s and the rising share of women in it could imply large differences across cohorts in the relative sizes of both groups. Because the community detection approach allows us to classify every author into one of the two groups and we observe the first publication of each of them, we can plot the type distribution across cohorts. Figure 8 does just that, plotting for each 5-year cohort of authors, the share assigned to the Osborne group under our benchmark specification. We find the Osborne group share to be stable around 39 percent for the 70s to 80s cohorts. The Osborne share grows during the 90s and up to the 2000-2004 cohort to around 46 percent, and stabilizes at around that number among the subsequent cohorts. Thus, although there is some compositional change over time towards the more liberal type, it is relatively modest and certainly insufficient to account on its own for the large changes in writing styles observed over the same period. The community assignment also allows us to compare both groups in terms of their observables. Figure A.14 plots the distributions of the main author-level characteristics we observe conditional on type. Across the board, both communities look very similar in their sex, ethnicity, and fields compositions, and have very similar productivity and citations distributions.

The reader may notice that we did not include a pairwise covariate capturing shared institutional affiliations. This is because we were unable to collect that information systematically across all authors in the network. A concern may be, thus, that the community assignment we estimated is mostly capturing shared institution-based homophily. To assuage this concern, we collected detailed institutional affiliation information for the subset of professors from the top-ranked 39 Economics programs, starting in 1990. For this sub-sample of authors we explored whether university affiliations predict community membership. We do so by running linear regressions of the 'Osborne-group' dummy variable on a dummy variable capturing affiliation, separately for each institution. Figure 9 presents a scatter plot of the coefficient on the university dummy against its corresponding p-value across all 39 regressions. Only 3 of them yield statistically significant coefficients at the 5 percent level, and 34 of them are smaller than 0.1 in magnitude. At least among this elite set of economic theorists, university affiliation does not predict the community assignment.

[28] That we always classify Osborne and Rubinstein in different groups is consistent their own expressed views we mentioned in the introduction.

6 Estimation of the writing-style model

6.1 Estimation

Armed with our community assignment from subsection 5.3, we set $O_i = 1$ for all authors classified in Osborne's community. Together with the control function estimates from subsection 4.2, we can write $u^{\rho}_{a(ij)t}$ from (1) as $u_{a(ij)t} = V^{\rho}_{a(ij)t} + \nu_{a(ij)t}$, where $V^{\rho}_{a(ij)t}(\beta_{i|\psi_i}, \beta_{j|\psi_j}) \equiv$

$$\widetilde{\alpha}^{\rho} + \varphi^{\rho}_t + \phi(\mathbf{z}_{ij})[\beta_{i|\psi_i} r^{\rho}_{it} + \delta^{\rho} O_i] + (1 - \phi(\mathbf{z}_{ij}))[\beta_{j|\psi_j} r^{\rho}_{jt} + \delta^{\rho} O_j] + \lambda^{\rho}[\hat{\eta}^{\rho}_{it} + \hat{\eta}^{\rho}_{jt}], \tag{7}$$

$\widetilde{\alpha}^{\rho} \equiv \alpha^{\rho} + \delta^{\rho}_R$, $\delta^{\rho} \equiv \delta^{\rho}_O - \delta^{\rho}_R$, and $\nu_{a(ij)t}$ is independent of $(r^{\rho}_{it}, r^{\rho}_{jt}, O_i, O_j)$ and type-1 extreme value distributed. As a functional form for the bargaining weights we use

$$\phi(\mathbf{z}_{ij}) = \frac{1}{1 + \exp\left(-\boldsymbol{\mu}' \mathbf{z}_{ij}\right)}$$

Under (7), and collecting in vector $\boldsymbol{\theta}$ all parameters, the unconditional likelihood of observing writing style $p_{a(ij)t} = \rho$ for paper $a(ij)t$ averages over the distribution of peer effects for each author conditional on their psychological types (conformist or contrarian), and then averages over the distribution of psychological types conditional on a vector of characteristics $\mathbf{w}_i$:

$$\mathbb{P}(p_{a(ij)t} = \rho | \mathbf{w}_i, \mathbf{w}_j, \mathbf{z}_{ij}; \boldsymbol{\theta}) =$$
$$\sum_{\psi_i, \psi_j \in \{\underline{\psi}, \overline{\psi}\}} \left[\int \int \frac{\exp\left(V^{\rho}_{a(ij)t}(\beta_{i|\psi_i}, \beta_{j|\psi_j})\right)}{1 + \sum_{s \in \{m,f,x\}} \exp\left(V^{s}_{a(ij)t}(\beta_{i|\psi_i}, \beta_{j|\psi_j})\right)} dF(\beta_i|\psi_i) dF(\beta_j|\psi_j) \right] \mathbb{P}(\psi_i|\mathbf{w}_i) \mathbb{P}(\psi_j|\mathbf{w}_j).$$

We think of conformism and contrarianism as psychological traits that are possibly stationary in the overall population. Over the last fifty years, however, the Economics profession has grown in size. For example, while we see 1,620 economists from the 1970s cohort, we see 4,970 from the 1990s cohort, and 11,317 from the 2010 cohort. The profession also has shifted its sex composition towards women. Because the new entrants or women as a whole could differ in their psychological inclinations relative to incumbents, we allow $\mathbf{w}_i = (\text{sex}_i, O_i)$ to include the authors' sex and community assignment.

The likelihood for the writing style choices across all articles is thus

$$L(\boldsymbol{\theta}|\mathbf{P}, \mathbf{W}, \mathbf{Z}) = \prod_{a} \prod_{\rho \in \{m,f,x,p\}} \mathbb{P}(p_{a(ij)t} = \rho | \mathbf{w}_i, \mathbf{w}_j, \mathbf{z}_{ij})^{\mathbf{1}\{p_{a(ij)t} = \rho\}}. \tag{8}$$

We use maximum simulated likelihood to estimate $\boldsymbol{\theta}$.[29] The vector of parameters to estimate includes the three pronoun specific intercepts $\widetilde{\alpha}^{\rho}$, the three sets of time effects φ_t^{ρ} (in practice we include time effects for groups of 5 years with the exception of 1970-1974 and 1975-1979 for which we include a single time effect[30]), the three coefficients δ^{ρ} on the Osborne community dummy, the three coefficients λ^{ρ} on the control function, the five coefficients on the pairwise covariates $\boldsymbol{\mu}$ on the bargaining weight function, the six location and scale parameters $(a_{\underline{\psi}}, b_{\underline{\psi}}, S_{\underline{\psi}}, a_{\overline{\psi}}, b_{\overline{\psi}}, S_{\overline{\psi}})$ governing the distribution of peer effect heterogeneity among conformists and contrarians, and the four conditional probabilities governing the psychological type distribution, $\mathbb{P}(\psi_i = \underline{\psi}|\mathbf{w})$.

6.2 Results

Our main estimation results for the writing-style choice model use the $n = 10$ acquaintance sets definition, and consider both past co-authoring and citation relations as channels of peer effects. We present the estimates in Table 9 and in Figure 10. The parameter estimates are all relative to the baseline choice, only plural. The year-group effects, φ_t^{ρ}, capture the overall trends in pronoun form popularity from Figure 1, with the masculine form trending down over time, the mixed and feminine forms trending up over time, and all forms converging to similar base rates near the end of the sample period. Notice also that the coefficients on the control function regressor, λ^{ρ}, are highly statistically significant, corroborating the importance of controlling for the endogeous component of the variation in peer influences. The pattern of estimates on δ^{ρ} reveal that relative to the plural form, on average over the 1970-2019 period authors in the Osborne community are considerably less likely to choose the masculine form (-0.61), somewhat less likely to choose the mixed form (-0.44), and almost equally likely to choose the feminine form (-0.08). These coefficients are all highly statistically significant. This pattern confirms to us that the subset of authors we identified as being part of Osborne's community does share affinity with his values related to the expression of views on gender in their writing. Moreover, it confirms that this otherwise unobserved dimension of preferences driving homophily in co-authoring is indeed related to relatively more or less traditional views. Turning to the co-authoring bargaining weights θ, we find that larger age and citations gaps do favor the most senior and cited author, while the coefficients for other pairwise characteristics are statistically insignificant.

Figure 10 presents the estimated distributions of peer effect heterogeneity among conformists (panel (a)) and among contrarians (panel (b)). While the average peer effects are

[29] See Appendix subsection 11.7 for additional details about the estimator.

[30] In 1970-1974 no papers used the only feminine choice, so the time effects for that group of years would be unidentified.

similar in magnitude, 11.3 among conformists and −9.2 among contrarians, there is much more heterogeneity among conformists. In fact, there is very little heterogeneity in peer effects among contrarians. The parameters of these Beta distributions are precisely estimated. At the mean of the data and all else equal, an increase in the masculine-only share of papers among his peers from 20 to 30 percent would raise the mean conformist's probability of writing an only-masculine paper from 12 to 38 percent. The partial effect would be of similar magnitude but of opposite direction for the mean contrarian. The implied partial effects are somewhat larger in magnitude for the feminine and mixed choices because their intercepts are less negative.

The last component of our model is the joint distribution of conformists and contrarians in the professional network. We present the estimated probabilities in panel (b) of Table 9. The majority of economists, 82 percent, are conformists, while only 18 percent are contrarians. Although there are some differences between men and women and between the more liberal (Osborne) community and the more conservative (Rubinstein) community, these are small. For example, Osborne-community women have the highest share of conformists, 90 percent. Rubinstenin-community men have the lowest share of conformists, 79 percent. These probabilities are very precisely estimated so the differences between these groups, albeit small, are statistically significant. An implication of these estimates is that the growth of the relatively more liberal group (see Figure 8), and the increasing participation of women (see Figure 2) have made the profession more conformist over time.

6.2.1 Model fit and robustness

Figure 11 presents a plot of the aggregate distribution of writing style choices over time from a simulation of our model using the estimated parameters. The exercise assumes starting values for the peer influences r_{i0}^{ρ} equal to the observed shares in 1970-1974, and simulates the choices going forward period to period holding fixed the set of papers and co-authorships in the data. While the model captures the downward trend in the masculine pronoun form, there is more persistence compared to the data, particularly early on. This is mostly because in the simulated data, the plural form does not show the increase in the 70s and 80s that it does in the data. The patterns of timing and ranking of the feminine and mixed choices, however, are very similar to the ones observed.

6.2.2 Choice-specific unobservables

We have argued that to first order, the value of the choices in our writing-style model is purely social, obviating the need for choice-specific unobservables. In this section we explore

several avenues of possible 'fundamental' differences in the value of choosing some writing styles over others.

Beliefs about journal editors' preferences. If authors believe that journal editors prefer some writing styles over others and make editorial decisions accordingly, this will induce differences in the perceived value of alternative choices. One possibility is that authors pay attention to the sex composition of the editorial boards of journals, and make writing style choices accordingly if, for example, they believe male and female editors have different preferences over writing styles and act on them.

To test for this possibility, we collected information on the editorial board membership of the "top five" general interest journals and five other highly ranked theory journals, from 1970 to the present. We then estimated linear probability models for the choice of only-masculine pronoun form at the article level on the average number of women in the editorial board in the three years prior to a paper's publication. Naturally, in this exercise we can only include the subset of papers published in any of the ten journals for which we have editorial information. We report the results in Table 6.

Column 2 reports estimates from a model that includes author fixed effects, using only within author variation across publications. Overall, the share of female editors does not predict gendered pronoun choice. In the remaining columns we explore whether there are differences between papers authored by men and women. We find no effects for men (columns 5 and 6), but large and positive effects for women, even in the specification that includes author fixed effects. Women are more likely to use a masculine only writing style when facing a larger share of female editors.[31]

Expectations of conformity by un-tenured professors. If writing styles are perceived to matter for publication and more broadly for career concerns, un-tenured economists may respond to such perceptions through their choices. The profession may expect, or may be perceived to expect, for example, more traditionalist attachments by younger economists. We test this possibility estimating linear probability models of gendered pronoun choice on a dummy variable equal to 1 for articles with at least an author in the first six years of his academic career. The top row of Table 7 presents the results. Authors at an early stage of their career are more likely to chose the plural writing style (relative to all other three

[31]This result may be surprising. Kosnik (2022), however, reports writing behavior among female economists that may be consistent with it. She studies writing style as measured by text sentiment in Economics articles published in prestigious journals, and finds that sentiment is more negative in papers written by women compared to men. She argues this is driven by career concerns, because papers with more negative writing styles tend to receive more citations.

styles). This is the case in both the specifications without (col. 5) and with (col. 6) author fixed effects.

Differential changes women's preferences We did not find significant differences between men and women in their preferences or their psychological types. Research in other areas has found, however, that women may effectively express their preferences only once they constitute a large enough share of the relevant network (e.g., Owen and Temesvary (2018) in the case of bank boards). A possibility in our setting is that the small share of female economic theorists, particularly early on, has limited their ability to signal their preferences more strongly. Because the profession has seen a steady growth in the share of women, in Table A.12 we explore this possibility by estimating differences in the distribution of the Osborne dummy for women of different cohorts. We do not detect any significant differences across the cohorts.

Varying signaling preferences across degrees of journal prestige. Could authors believe some writing styles to be more or less appropriate at journals of different degrees of professional prestige and act accordingly? We test this possibility in Table 7 estimating linear probability models of gendered pronoun choice on either the log ranking of the article's journal (second row), or a dummy variable equal to 1 for articles in either of the "top five" journals in the profession. Articles in more prestigious journals are indeed less likely to use plural forms (col. 6), and more likely to use mixed forms (col. 8), even after including author fixed effects.

Underlying complementarities between sub-fields and writing styles. Could it be that authors' writing style choices respond to the paper's topic, as captured by it's sub-field within Economic Theory? For example, contract theory could be more amenable to the mixed writing style if the principal-agent dichotomy is projected onto the masculine-feminine binary. Or perhaps more abstract sub-fields may be more amenable to the plural form. We assess this possibility in Table 8, where we present estimates of linear probability models of the different writing styles on theory sub-field dummies. Even-numbered columns present estimates from models that include author fixed effects, and reveal that authors are less likely to chose plural forms and more likely to chose mixed forms, when publishing in Collective decision-making, Game theory, Information economics, and Welfare economics.

Motivated by these findings, we estimate a specification of our writing-style model where we include the following variables as additional shifters of the choice-specific payoffs: i) a dummy variable equal to 1 if the article has at least one female author, ii) a dummy variable

equal to 1 if the article has at last one un-tenured author, iii) the journal's ranking, and iv) a dummy variable for whether the article's sub-field is either Collective decision-making, Game theory, Information economics, or Welfare economics.

6.2.3 Decomposing the roles of cohorts, women, co-authorship and peer influence

In this section we use the parameter estimates and our model to assess the quantitative importance of the different margins contributing to the observed change in behavior over the last fifty years. We do this taking the set of articles and the professional network relations fixed, and simulating pronoun form choices for each paper, and computing the resulting peer influence variables r_{it}^{ρ} for each year that then determine choices for subsequent articles. As an initial condition for these simulations we take the average 1970-1974 observed average distribution of choices. To highlight the roles played by the different components, we undertake these simulations in a 'stationary' environment where we zero-out the estimated time trends φ_t^{ρ}.

The first column of Table 10 presents the average (2014-2019) distribution of pronoun form choices in this stationary baseline scenario. In all other columns we report percentage point differences relative to the baseline simulation. In the fifth column we quantify the importance of social interactions, by reporting the corresponding distribution of choices in a simulation where we shut down all peer influences by making $\beta_i = 0$ for all authors. Of course, in this case only author-level idiosyncratic preferences and the patterns of co-authorship matter. Compared to the baseline scenario, the absence of peer effects would lead to a 20 percentage points higher feminine only share of papers, a 19 percentage points higher mixed share of papers, a 28 percentage points higher share of plural form papers, and a 68 percentage points lower share of masculine only papers. This exercise highlights that the large fraction of economists exhibiting positive peer effects has been a major drag to the faster diffusion of the more innovative writing styles, given the initially overwhelming prevalence of the masculine form.

To assess the importance of co-authoring, the last column of Table 10 instead leaves peer effects as in the baseline, but instead asks how the long-term evolution would have differed if the share of co-authored articles had remained as in the early 70s (65 percent) throughout the whole period. We implement this exercise by altering, for a subset of co-authored articles chosen at random (at the rates that make the year by year distribution of publications match the 1970s co-authoring rate), the identity of one of the co-authors, turning the article into a single authored one. This exercise makes no difference for the 2014-2019 aggregate distribution of choices, suggesting that increased co-authorship was not

a major contributor to the increasing variability in observed choices. This is, to a large extent, because co-authorship is highly homophilous in the theorists' professional network.

In the second to fourth columns of Table 10 and in Figure 12 we instead explore how the observed distribution of psychological profiles contributes to the long-term patterns in writing style. We do this by reporting the 2014-2019 average distribution of choices, in simulations where we vary the underlying distribution of psychological types, from 100 percent contrarians to 100 percent conformists. The vertical dashed line is located at the estimated distribution. The figure reveals a non-monotonic relation between the share of conformists and the distribution of pronoun forms: in the range in which conformists are in the minority, the share of feminine and mixed forms grows as the conformist share increases, at the expense of the plural and the masculine. Notice that at the author level, there is a considerable amount of switching driven by the large numbers of contrarians, although at the aggregate level the shares are smooth functions of the share of conformists. Around when the distribution of conformists and contrarians is balanced, when conformists become the majority the masculine share starts growing at the expense of the feminine and mixed forms. These two fall in popularity as more conformists appear in the population. In fact, the plot reveals that around 50 percent conformists there is a tipping point: the contrarian share is too small to prevent the positive feedback forces from the large share of conformists taking over, and the masculine form grows dominant.

7 Conclusions

We study the transformation in writing styles within the academic community of economic theorists between 1970 and 2019. In this period, the choice of gender for third person pronouns in the publications of these academics has moved away from the exclusive use of the masculine form to the adoption of plural and feminine pronoun forms, revealing changing views about gender more broadly. During the same period, the profession saw a large expansion in its size, increased academic collaboration, and an increasing participation of women. Against this background, we use a discrete choice model of writing style to quantify the importance of peer influences within the professional network in driving the long term changes we observe. As a source of exclusion restrictions to identify the peer effects, we build an underlying network of feasible connections in the professional network borrowing recent tools from the Natural Language Processing literature. We highlight that peer influences are of two kinds: while some economists are conformists (they move towards their peers' choices), others are contrarians (they move away from their peers' choices). The presence of these competing behaviors is a key driver of the dynamics of adoption of the

new pronoun forms. Our model allows us to quantify their importance, and to decompose the changes in writing styles between cohort effects, the entry of women to the profession, increased co-authoring, and peer influence.

References

Akerloff, G. and R. Kranton (2000): "Economics and Identity," *Quarterly Journal of Economics*, 115, 715–753.

Akerloff, G., J. Yellen, and M. Katz (1996): "An Analysis of Out-of-Wedlock Childbearing in the United States," *Quarterly Journal of Economics*, 111, 469–530.

Alesina, A., P. Giuliano, and N. Nunn (2013): "On the Origins of Gender Roles: Women and the Plough," *Quarterly Journal of Economics*, 128, 469–530.

Anderson, K. A. and S. Richards-Shubik (2022): "Collaborative Production in Science: An Empirical Analysis of Coauthorships in Economics," *The Review of Economics and Statistics*, 104, 1241–1255.

Angrist, J., G. Imbens, and D. Rubin (1996): "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, 91, 444–472.

Baron, D. (1986): *Grammar and Gender*, Yale University Press.

Bayer, P. and C. Timmins (2007): "Estimating Equilibrium Models of Sorting Across Locations," *The Economic Journal*, 117, 353–374.

Besancenot, D., K. Huynh, and F. Serranito (2017): "Co-authorship and research productivity in economics: Assessing the assortative matching hypothesis," *Economic Modelling*, 66, 61–80.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.

Bramoulle, Y., H. Djebbari, and B. Fortin (2009): "Identification of Peer Effects through Social Networks," *The Journal of Econometrics*, 150, 41–55.

Chamberlain, G. (1980): "Analysis of Covariance with Qualitative Data," *Review of Economic Studies*, 47, 225–238.

Chari, A. and P. Goldsmith-Pinkham (2017): "Gender Representation in Economics across Topics and Time: Evidence from the NBER Summer Institute," NBER Working Paper No. 23953.

Combes, S. and P. Givord (2018): "Selective Matching: Gender Gap and Network Formation in Research," Crest Working Paper No. 2018-07.

Dahl, C., M. Huber, and G. Mellace (2023): "It is Never too LATE: A New Look at Local Average Treatment Effects with or without Defiers," *The Econometrics Journal*, 26, 378–404.

dePaula, A., S. Richards-Shubik, and E. Tamer (2018): "Identifying Preferences in Networks with Bounded Degree," *Econometrica*, 86, 263–288.

Ductor, L., S. Goyal, and A. Prummer (2023): "Gender and Collaboration," *The Review of Economics and Statistics*, 105, 1366–1378.

Ductor, L. and A. Prummer (2023): "Gender Homophily, Collaboration, and Output," Working Paper, Universidad de Granada.

Fafchamps, M., S. Goyal, and marco van der Leij (2010): "Matching and Network Effects," *Journal of the European Economic Association*, 8, 203–231.

Feng, Y., S. Huang, and J. Sun (2023): "Pairwise Covariates-Adjusted Block Model for Community Detection," ArXiv.

Fernández, R., S. Parsa, and M. Viarengo (2024): "Coming out in America: thirty years of cultural change," *The Journal of Law, Economics, and Organization*, ewae010.

Freeman, R. B. and W. Huang (2014): "Collaborating with People like me: Ethnic co-authorship within the US," NBER Working Paper No. 19905.

Fryer, R. and S. Levitt (2004): "The Causes and Consequences of Distinctively Black Names," *Quarterly Journal of Economics*, 119, 767–805.

García-Jimeno, C., A. Iglesias, and P. Yildirim (2022): "Information Networks and Collective Action: Evidence from the Women's Temperance Crusade," *American Economic Review*, 112, 41–80.

Giuliano, P. and N. Nunn (2021): "Understanding Cultural Persistence and Change," *Review of Economic Studies*, 88, 1541–1581.

Goldin, C. (2006): "The Quiet Revolution that Transformed Women's Employment, Education, and Family," *American Economic Association Papers and Proceedings*, 96, 1–21.

——— (2023): "Why Women Won," NBER Working Paper No. 31762.

Goldin, C. and M. Shim (2004): "Making a name: Women's Surnames at Marriage and Beyond," *Journal of Economic Perspectives*, 18, 143–160.

Goldin, C. D. (1991): "The role of World War II in the rise of women's employment," *The American Economic Review*, 741–756.

Goolsbee, A. and P. Klenow (2002): "Evidence on Learning and Network Externalities in the Diffusion of Home Computers," *Journal of Law and Economics*, 45, 317–343.

Goyal, S., M. van del Leij, and J. L. Moraga-Gonzalez (2006): "Identification of Peer Effects through Social Networks," *The Journal of Political Economy*, 114, 403–412.

Graham, B. (2017): "An Econometric Model of Network Formation with Degree Heterogeneity," *Econometrica*, 85, 1033–1063.

Griliches, Z. (1957): "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica*, 25, 501–522.

Grosjean, P. and R. Khattar (2019): "It's raining men! Hallelujah? The long-run consequences of male-biased sex ratios," *The Review of Economic Studies*, 86, 723–754.

Hammermesh, D. (2013): "Six Decades of Top Economics Publishing: Who and How?" *Journal of Economic Literature*, 51, 162–172.

——— (2015): "Age, Cohort, and Coauthorship," NBER Working Paper No. 20938.

Islam, A., J. S. Kim, S. Song, and Y. Zenou (2022): "Network Formation with Unobserved Homophily: Identification and Consistent Estimation," Monash University.

Iyengar, R., C. V. den Bulte, and T. Valente (2011): "Opinion Leadership and Social Contagion in New Product Diffusion," *Marketing Science*, 30, 195–212.

Jochmans, K. (2023): "Peer Effects and Endogenous Social Interactions," *Journal of Econometrics*, 235, 1203–1214.

Johnsson, I. and H. R. Moon (2021): "Estimation of Peer Effects in Endogenous Social Networks: Control Function Approach," *The Review of Economics and Statistics*, 103, 328–345.

Kandori, M. (1992): "Social norms and community enforcement," *Review of Economic Studies*, 59, 63–80.

Karni, E. and D. Schmeidler (1990): "Fixed Preferences and Changing Tastes," *American Economic Association Papers and Proceedings*, 80, 262–267.

Karrer, B. and M. Newman (2010): "Stochastic Block Models and Community Structure in Networks," ArXiv.

Kelli, M. and C. O. Gráda (2000): "Market Contagion: Evidence from the Panics of 1854 and 1857," *American Economic Review*, 90, 1110–1124.

Kosnik, L.-R. (2022): "Who Are the More Dismal Economists? Gender and Language in Academic Economics Research," *American Economic Association Papers and Proceedings*, 112, 1–5.

Kuld, L. and J. O'Hagan (2018): "Rise of multi-authored papers in economics: Demise of the 'lone star' and why?" *Scientometrics*, 114, 1207–1225.

Lieberson, S. and E. Bell (1992): "Children's First Names: An Empirical Study of Social Taste," *American Journal of Sociology*, 98, 511–554.

Lindenlaub, I. and A. Prummer (2020): "Network Structure and Performance," *Economic Journal*, 131, 851–898.

Lundberg, S. and J. Stearns (2019): "Women in economics: Stalled progress," *Journal of Economic Perspectives*, 33, 3–22.

Manheim, K. (1952): *Essays on the Sociology of Knowledge*, Oxford, UK: Oxford U Press.

Matsuyama, K. (1991): "Custom versus Fashion: Hysteresis and Limit Cycles in a Random Matching Game," Working Paper, Northwestern U.

McDowell, J. M. and M. Melvin (1983): "The Determinants of Co-Authorship: An Analysis of the Economics Literature," *The Review of Economics and Statistics*, 65, 155–160.

Mikolov, T., K. Chen, G. Corrado, and J. Dean (2013): "Efficient Estimation of Word Representations in Vector Space," ArXiv Working Paper 1301.3781v3.

Mullahy, J. (2011): "Multivariate fractional regression estimation of econometric share models," Working Paper 16354, University of Wisconsin-Madison.

Nevo, A. (2003): "Measuring Marker Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69, 307–342.

Newman, M. (2001): "The Structure of Scientific Collaboration Networks," *Proceedings of the National Academy of Sciences*, 98, 404–409.

——— (2018): *Networks*, Oxford, UK: Oxford U Press.

Onucich, P. and D. Ray (2021): "Signaling and Discrimination in Collaborative Projects," *American Economic Review*, 113, 210–252.

Osborne, M. and A. Rubinstein (1994): *A Course in Game Theory*, Cambridge, MA: MIT Press.

Owen, A. and J. Temesvary (2018): "The Performance Effects of Gender Diversity on Bank Boards," *Journal of Banking and Finance*, 90, 50–63.

Rayder, N. (1965): "The Cohort as a Concept in the Study of Social Change," *American Sociological Review*, 30, 843–861.

Sarsons, H., K. Gërxhani, E. Reuben, and A. Schram (2021): "Gender Differences in Recognition for Group Work," *Journal of Political Economy*, 129, 101–147.

Stevenson, B. and H. Zlotnick (2018): "Representations of Men and Women in Introductory Economics Textbooks," *American Economic Association Papers and Proceedings*, 108, 180–185.

Young, P. (2014): "The Evolution of Social Norms," Working Paper No. 726, Oxford University.

Önder, A. S., S. Schweitzer, and H. Yilmazkuday (2021): "Specialization, field distance, and quality in economists' collaborations," *Journal of Informetrics*, 15, 1751–1577.

8 Tables

	(1) Co-authors	(2) Acquaintances	(3) Non-coauthors	(4) All
Same ethnicity	0.38	0.23	0.16	0.16
	(0.49)	(0.42)	(0.36)	(0.36)
Same sex	0.77	0.76	0.71	0.71
	(0.42)	(0.43)	(0.45)	(0.45)
Common fields	1.41	1.00	0.29	0.29
	(0.98)	(0.91)	(0.56)	(0.56)
Age difference	9.20	10.89	13.67	13.67
	(8.83)	(9.16)	(10.72)	(10.72)
Citations difference	4,720	5,060	2,457	2,457
	(12,576)	(12,335)	(7,244)	(7,244)
Productivity difference	12.51	11.91	5.07	5.07
	(15.99)	(14.92)	(8.65)	(8.65)
Log productivity product	3.53	3.42	1.70	1.70
	(1.94)	(1.71)	(1.38)	(1.38)
Pairs	50,778	748,023	429,238,173	429,288,951

Table 1: Pairwise Characteristics. The table reports means and standard deviations (in parenthesis) for a set of pairwise characteristics across pairs of economists in the professional network. Column (1) restricts the set to include only pairs of economist who co-authored together. Column (2) restricts the set to include only pairs of economists in each others acquaintance sets. Column (3) restricts the set to include only pairs of economists who never co-authored with each other. Column (4) includes all pairs of economists in the professional network.

Panel A	Transition matrix for all sequences of pairs of articles			
	Masculine	Feminine	Plural	Mixed
From/To	(1)	(2)	(3)	(4)
Masculine	0.52	0.06	0.24	0.18
Feminine	0.18	0.31	0.24	0.26
Plural	0.28	0.09	0.49	0.14
Mixed	0.28	0.14	0.20	0.38
Panel B	Implied stationary distributions			
	Masculine	Feminine	Plural	Mixed
	(1)	(2)	(3)	(4)
Overall	0.35	0.12	0.31	0.22
Only single-authored	0.43	0.09	0.29	0.19
Only 70s cohort	0.51	0.04	0.29	0.16
Only 80s cohort	0.39	0.08	0.33	0.20
Only 90s cohort	0.33	0.12	0.31	0.24
Only 00s cohort	0.29	0.16	0.30	0.25
Only 10s cohort	0.26	0.19	0.31	0.24

Table 2: Overall transition matrix and stationary distributions. Panel A presents the implied transition matrix across all sequential pairs of articles. Panel B presents the implied stationary distribution for the transition matrix in Panel A (overall), and the transition matrices that restrict attention to sequential single-authored pairs of articles, and for all sequential pairs of articles by author cohorts. The corresponding transition matrices for the single-authored and cohort groups appear in Table A.11.

	Fractional Multinomial Response Models		
	Social network: Co-authors and cited		
	Dep var: Share of articles by author i's social network using writing style		
	Masculine	Feminine	Mixed
	(1)	(2)	(3)
Δz_{it}^m	2.05	0.78	-1.44
	(0.04)	(0.06)	(0.05)
Δz_{it}^f	-0.96	5.20	1.91
	(0.06)	(0.10)	(0.07)
Δz_{it}^x	-1.55	2.42	3.91
	(0.06)	(0.08)	(0.08)
Obs.		68,944	
	Social network: Only co-authors		
	Dep var: Share of articles by author i's social network using writing style		
	Masculine	Feminine	Mixed
	(4)	(5)	(6)
Δz_{it}^m	1.72	0.57	-1.12
	(0.04)	(0.05)	(0.04)
Δz_{it}^f	-0.29	3.92	1.65
	(0.06)	(0.09)	(0.07)
Δz_{it}^x	-1.23	1.84	2.91
	(0.05)	(0.07)	(0.07)
Obs.		67,733	
	Social network: Only cited		
	Dep var: Share of articles by author i's social network using writing style		
	Masculine	Feminine	Mixed
	(7)	(8)	(9)
Δz_{it}^m	2.96	1.52	-2.55
	(0.11)	(0.13)	(0.11)
Δz_{it}^f	-2.44	5.72	1.52
	(0.14)	(0.21)	(0.15)
Δz_{it}^x	-2.46	4.52	7.13
	(0.13)	(0.18)	(0.18)
Obs.		68,903	

Table 3: Control Function Models of Pronoun Choice. The table presents coefficient estimates of the fractional multinomial choice conditional mean equations. The explanatory regressors measure the change in average pronoun choice of peers of a given author's peers who are not his acquaintances. The baseline category is the plural form. The top panel considers co-authors and citees as peers. The middle panel considers only co-authors as peers. The bottom panel considers only citees as peers.

	Social network: Co-authors and cited			
	Masculine	Feminine	Mixture	Plural
	(1)	(2)	(3)	(4)
z_{it}^m	0.42	-0.02	0.03	-0.44
	(0.01)	(0.00)	(0.01)	(0.01)
z_{it}^f	-0.15	0.50	0.08	-0.43
	(0.01)	(0.01)	(0.01)	(0.01)
z_{it}^x	-0.03	0.07	0.50	-0.55
	(0.01)	(0.01)	(0.01)	(0.01)
Authors FEs	Y	Y	Y	Y
R^2	0.76	0.67	0.75	0.72
F-statistic	228	390	214	69
	Social network: Only co-authors			
	Masculine	Feminine	Mixture	Plural
	(5)	(6)	(7)	(8)
z_{it}^m	0.38	-0.02	0.04	-0.40
	(0.01)	(0.00)	(0.01)	(0.01)
z_{it}^f	-0.19	0.48	0.10	-0.38
	(0.01)	(0.01)	(0.01)	(0.01)
z_{it}^x	-0.03	0.08	0.44	-0.50
	(0.01)	(0.01)	(0.01)	(0.01)
Authors FEs	Y	Y	Y	Y
R^2	0.65	0.59	0.62	0.62
F-statistic	102	209	94	59
	Social network: Only cited			
	Masculine	Feminine	Mixture	Plural
	(9)	(10)	(11)	(12)
z_{it}^m	0.56	0.03	0.08	-0.67
	(0.02)	(0.01)	(0.01)	(0.02)
z_{it}^f	0.33	0.40	-0.08	-0.66
	(0.03)	(0.01)	(0.02)	(0.03)
z_{it}^x	0.03	0.00	0.69	-0.73
	(0.03)	(0.01)	(0.02)	(0.03)
Authors FEs	Y	Y	Y	Y
R^2	0.80	0.72	0.87	0.80
F-statistic	480	376	513	44
Obs.	84,565	84,565	84,565	84,565

Table 4: Robustness: Linear Models for Pronoun Choice. The table presents coefficient estimates of the within-author panel linear regression models for the four pronoun form shares. The explanatory regressors measure the average pronoun choice of peers of a given author's peers who are not his acquaintances. The baseline category is the plural form. The top panel considers co-authors and citees as peers. The middle panel considers only co-authors as peers. The bottom panel considers only citees as peers.

Pairwise Covariate	Acquaintance Set Definition $Q_{10}(i)$	$Q_5(i)$	$Q_{20}(i)$
$\underline{\gamma}$			
Same ethnicity	1.03	0.92	1.13
	(0.10)	(0.08)	(0.12)
Same sex	0.24	0.23	0.25
	(0.12)	(0.10)	(0.14)
Common fields	0.76	0.67	0.85
	(0.05)	(0.04)	(0.06)
Age difference	-1.98	-1.64	-2.29
	(0.51)	(0.44)	(0.61)
Citations difference	-4.86	-4.65	-4.96
	(0.19)	(0.17)	(0.23)
Productivity difference	0.34	0.21	0.36
	(0.49)	(0.41)	(0.58)
Log Productivity Product	0.49	0.47	0.51
	(0.02)	(0.02)	(0.03)
$\underline{\Omega}$			
$\omega_{\ell\ell}$	0.18	0.34	0.09
	(0.01)	(0.02)	(0.01)
$\omega_{\ell c}$	0.02	0.05	0.01
	(0.01)	(0.01)	(0.01)
ω_{cc}	0.06	0.11	0.03
	(0.02)	(0.03)	(0.02)
Rubinstein-type share	0.56	0.55	0.57

Table 5: Community Detection Estimates. The table presents maximum likelihood estimates of the covariates-adjusted stochastic block model for community detection (Feng et al., 2023). The first column presents results under the ten-closest acquaintance set definition. The second column presents results under the five-closest acquaintance set definition. The third column presents results under the 20-closest acquaintance set definition. The model is estimated on the 29,302 authors who co-authored at least once.

	Dependent variable: Only masculine pronouns dummy					
	Overall		With female author(s)		With only male authors	
	(1)	(2)	(3)	(4)	(5)	(6)
Exposure	0.16	-0.07	0.75	1.59	0.08	-0.15
	(0.07)	(0.12)	(0.20)	(0.61)	(0.08)	(0.12)
Year FEs	Y	Y	Y	Y	Y	Y
Journal FEs	Y	Y	Y	Y	Y	Y
Author FEs	N	Y	N	Y	N	Y
Obs.	10,918	6,804	1,465	519	9,453	6,013

Table 6: Exposure to Female Editors. The table presents linear probability models at the article level estimated by OLS, on the sub-sample of articles published by the authors from our theorists professional network in one of ten major Economics journals (top-5 general interest and the top 5 economic theory journals based on *RePEC*'s rankings of September 2023). The exposure variable is the average number of female editors of in the board of a journal, over the three years period prior to an article's publication date. Besides year and journal fixed effects, odd-numbered columns also include the date of first publication, the total number of publications, the total number of citations, the ethnicity, and the community assignment (Osborne/Rubinstein) of each author. Columns (1) and (2) include all articles in any of the ten journals. Even-numbered columns include author fixed effects instead. Columns (3) and (4) only include articles with at least one female author. Columns (5) and (6) only include articles with both male authors.

	Masculine		Feminine		Plural		Mixed	
Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
First 6 Years	-0.022	-0.011	0.018	-0.005	0.022	0.023	-0.018	-0.007
	(0.004)	(0.007)	(0.002)	(0.004)	(0.004)	(0.007)	(0.003)	(0.006)
Log(Rank)	0.006	0.002	-0.001	0.001	0.016	0.008	-0.022	-0.011
	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)
Top 5 Journal	0.019	0.006	0.005	-0.002	-0.073	-0.020	0.049	0.016
	(0.006)	(0.008)	(0.003)	(0.004)	(0.006)	(0.007)	(0.005)	(0.007)
Year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Author FEs	N	Y	N	Y	N	Y	N	Y
Obs.	66,533	48,632	66,533	48,632	66,533	48,632	66,533	48,632

Table 7: Alternative Drivers of Writing Style Choices. The table presents coefficient estimates from linear probability models at the article level, separately regressing dummy variables for each type of writing style on three different variables. In the first row we report results for models that include a dummy variable equal to 1 if at least one author is, at the time of publishing the paper, at most 6 years since his first publication, as a proxy for the tenure track period. In the second row we report results for models that include the log rank of the journal where the article was published, based on the most recent ranking here: www.researchbite.com. It combines an h-index, an impact score, and the SJR score. In the third row we report results for models that include a dummy variable equal to 1 if the journal where the article was published is either *Econometrica*, *The Review of Economic Studies*, *The Journal of Political Economy*, *The American Economic Review*, or *The Quarterly Journal of Economics*. Odd columns present results for models without author fixed effects. Even columns present results for models with author fixed effects instead.

	Masculine		Feminine		Plural		Mixed	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Analysis of Collective Decision-Making	0.038	0.018	0.053	-0.001	-0.234	-0.067	0.144	0.050
	(0.011)	(0.017)	(0.009)	(0.012)	(0.010)	(0.014)	(0.011)	(0.017)
Distribution	0.001	0.006	0.044	0.023	-0.012	-0.030	-0.033	0.002
	(0.027)	(0.038)	(0.021)	(0.027)	(0.029)	(0.039)	(0.020)	(0.034)
Financial Economics	0.027	0.005	-0.012	0.001	-0.035	-0.004	0.020	-0.003
	(0.007)	(0.011)	(0.004)	(0.006)	(0.007)	(0.010)	(0.006)	(0.009)
Game Theory	0.072	-0.012	0.037	-0.005	-0.232	-0.030	0.123	0.047
	(0.008)	(0.013)	(0.006)	(0.009)	(0.007)	(0.011)	(0.008)	(0.012)
General Equilibrium	0.090	0.022	0.019	0.006	-0.087	-0.011	-0.023	-0.018
	(0.016)	(0.021)	(0.011)	(0.014)	(0.016)	(0.020)	(0.012)	(0.017)
Household Behavior and Family Economics	-0.053	-0.057	0.016	-0.046	-0.096	-0.042	0.133	0.145
	(0.027)	(0.046)	(0.023)	(0.035)	(0.031)	(0.045)	(0.030)	(0.044)
Information, Knowledge, and Uncertainty	0.052	0.001	0.015	-0.003	-0.239	-0.064	0.172	0.066
	(0.009)	(0.014)	(0.007)	(0.010)	(0.008)	(0.011)	(0.009)	(0.014)
Market Structure, Pricing, and Design	-0.005	0.012	0.012	-0.004	-0.062	-0.010	0.055	0.001
	(0.008)	(0.013)	(0.006)	(0.009)	(0.009)	(0.013)	(0.007)	(0.012)
Micro-Based Behavioral Economics	-0.032	0.010	0.046	0.010	-0.147	-0.052	0.132	0.032
	(0.023)	(0.041)	(0.020)	(0.032)	(0.024)	(0.034)	(0.025)	(0.046)
Production and Organizations	-0.020	0.055	-0.032	-0.018	-0.040	-0.038	0.091	0.002
	(0.017)	(0.027)	(0.011)	(0.015)	(0.018)	(0.024)	(0.016)	(0.025)
Welfare Economics	0.026	-0.003	0.067	0.021	-0.177	-0.048	0.084	0.030
	(0.012)	(0.017)	(0.010)	(0.012)	(0.011)	(0.016)	(0.012)	(0.016)
Year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Author FEs	N	Y	N	Y	N	Y	N	Y
Obs.	66,533	48,637	66,533	48,637	66,533	48,637	66,533	48,637

Table 8: Differences in pronoun style choice by sub-fields.

Panel A: Parameters

	Masculine	Feminine	Mixed
	(1)	(2)	(3)
$\widetilde{\alpha}^{\rho}$ (Intercepts)	-1.65	-0.49	-0.35
	(0.08)	(0.04)	(0.05)
δ^{ρ} (Community dummies)	-0.61	-0.08	-0.44
	(0.07)	(0.04)	(0.04)
λ^{ρ} (Control functions)	-2.12	2.71	2.77
	(0.24)	(0.31)	(0.23)
φ^{ρ} (Year-group dummies)			
1970-1979	3.17	-4.57	-1.54
	(0.12)	(0.71)	(0.16)
1980-1984	2.68	-3.86	-0.73
	(0.13)	(0.45)	(0.11)
1985-1989	2.07	-2.60	-0.20
	(0.11)	(0.17)	(0.07)
1990-1994	1.46	-1.44	0.23
	(0.10)	(0.09)	(0.06)
1995-1999	0.74	-0.63	0.39
	(0.08)	(0.06)	(0.06)
2000-2004	0.93	-0.34	0.62
	(0.07)	(0.05)	(0.05)
2005-2009	1.00	-0.02	0.58
	(0.07)	(0.05)	(0.05)
2010-2014	1.09	0.07	0.53
	(0.07)	(0.05)	(0.05)
ϕ (Bargaining power)			
Age difference		0.98	
		(0.54)	
Citations difference		1.07	
		(0.67)	
Productivity difference		0.42	
		(0.51)	
Sex difference		0.08	
		(0.11)	
Shared ethnicity		-0.27	
		(0.08)	

Panel B: Distribution of types

Community	Sex	$\mathbb{P}$(Conformist|.,.)	$\mathbb{P}$(Contrarian|.,.)
	Men	0.85	0.15
Osborne		(0.02)	
	Women	0.90	0.10
		(0.03)	
	Men	0.79	0.21
Rubinstein		(0.02)	
	Women	0.82	0.18
		(0.03)	
Observations			56,473

Table 9: Parameter Estimates of the Writing Style Model. Both co-authoring and citations networks.

	Percentage points difference relative to baseline share by 2014-2019					
Choice	Baseline Share	100% Contrarian	100% Conformist	50% Contrarian	No Peer Effects	Baseline Coauthorship
Feminine	0.07	0.15	-0.07	0.20	0.20	0.00
Masculine	0.75	-0.67	0.25	-0.68	-0.67	0.01
Mixed	0.07	0.15	-0.07	0.19	0.19	0.00
Plural	0.12	0.37	-0.12	0.29	0.28	0.00

Table 10: Simulated end-line writing style shares under alternative scenarios. The table presents the difference in the average distributions of pronoun form shares in 2014-2019 relative the the baseline simulation in a stationary environment (zeroing out the estimated time effects) using the remaining estimated parameters from Table 9, under alternative scenarios. The first column presents the end point distribution under the baseline stationary simulation. The second column supposes that all authors are contrarians. The third column instead supposes that all authors are conformists. The fourth column supposes the distribution of conformists and contrarians is even. The fifth column supposes that there are no peer effects for any author. The last column supposes that the co-authorship rate remains at its average 1970-1974 value. In all simulations, the observed co-authorships are held fixed.

9 Figures

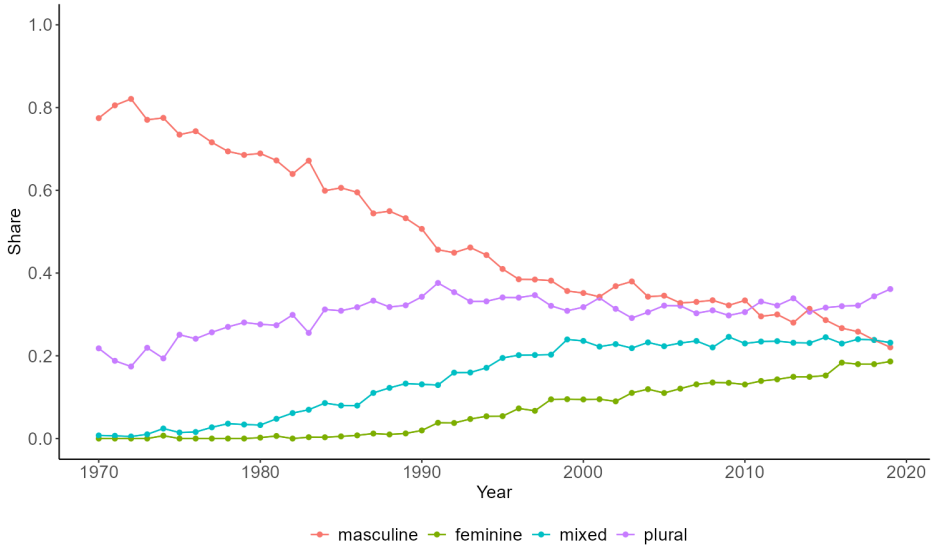
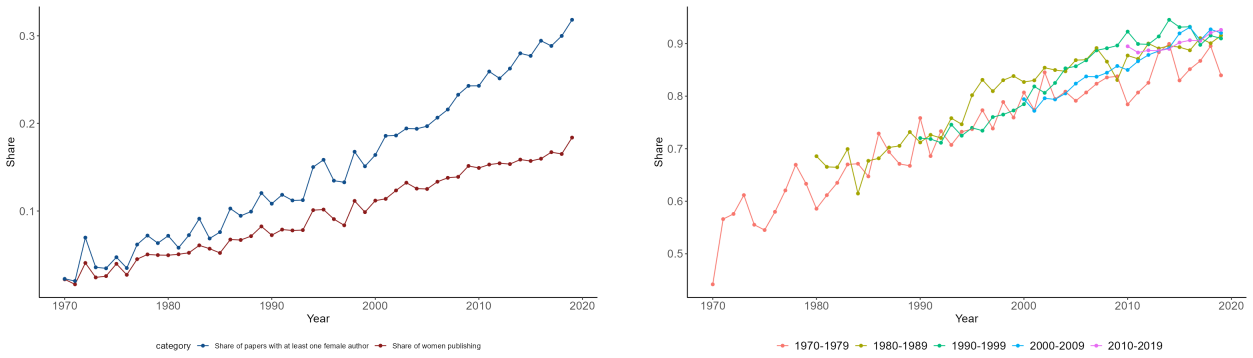


Figure 1: Distribution of pronoun use over time in theory papers, 1970-2019.



(a) Participation of women in the economic theory.

(b) Share of co-authored papers, by cohort.

Figure 2: Long-term change in the economics profession.

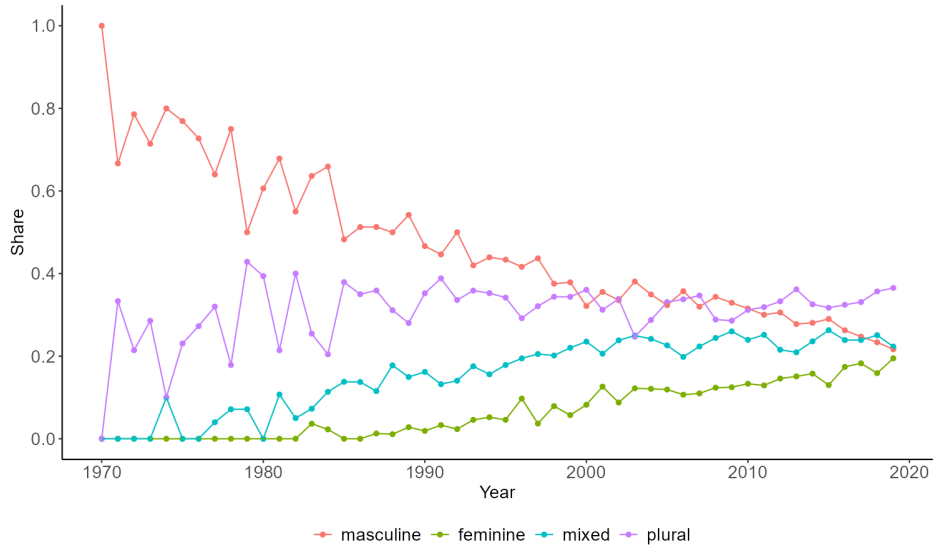
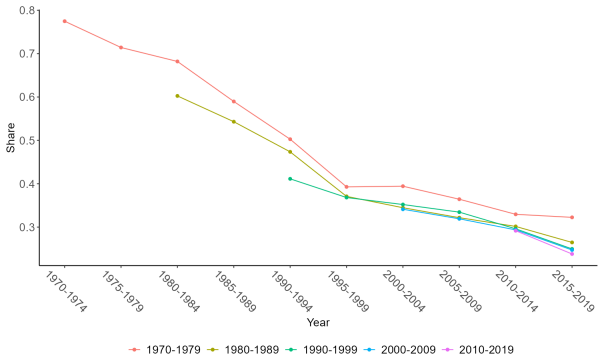
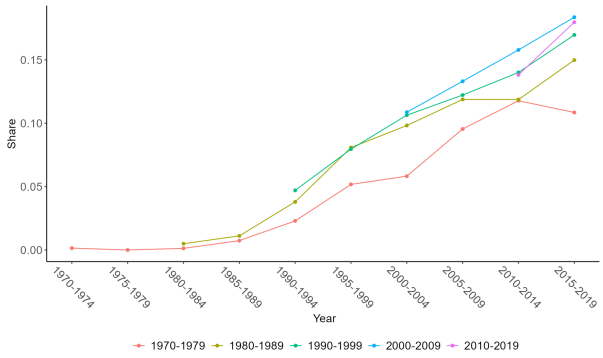


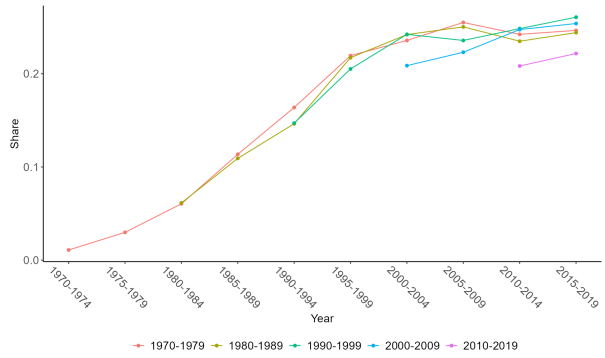
Figure 3: Distribution of pronoun use over time for paper authored by women, 1970-2019.



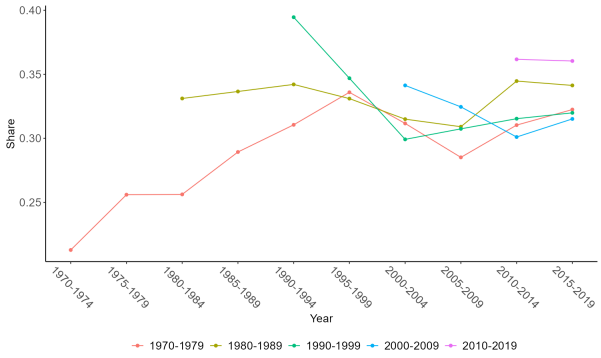
(a) Share of authors using masculine only.



(b) Share of authors using feminine only.

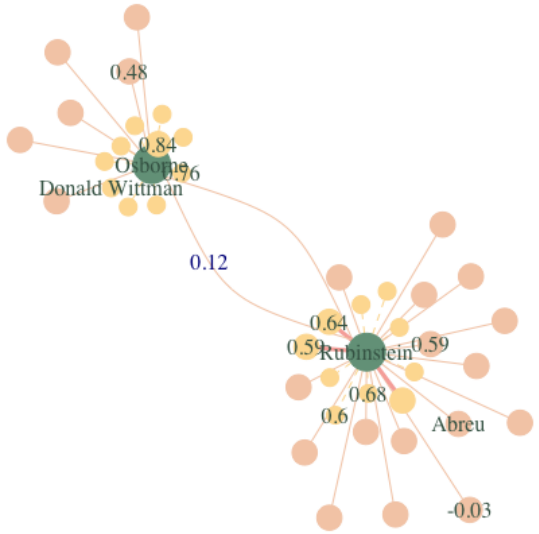


(c) Share of authors using mixed.

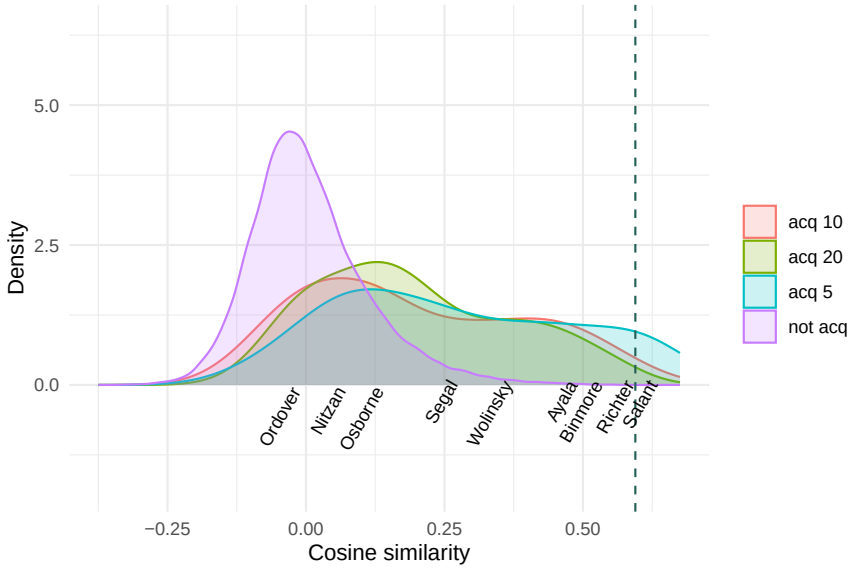


(d) Share of authors using plural only.

Figure 4: Distribution of pronoun use over time, by cohorts.

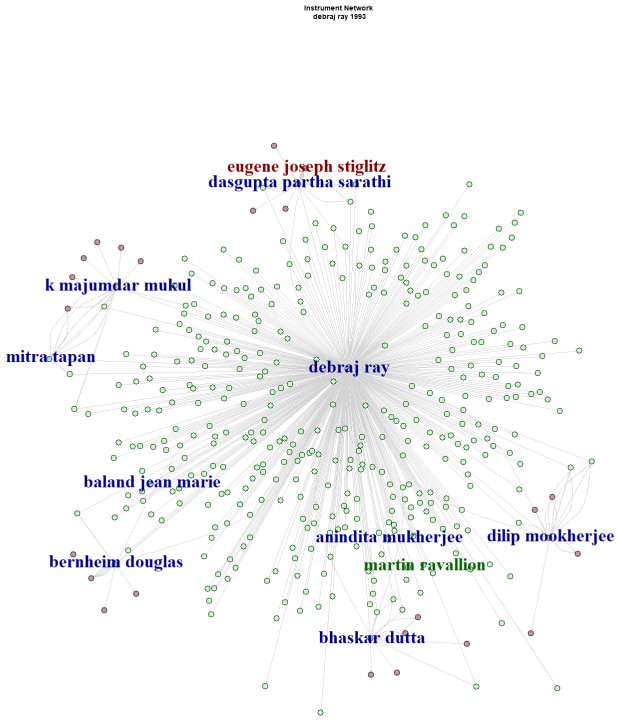


(a) Osborne and Rubinstein's local co-author network. Solid edges represent co-authorships. Dashed edges represent acquaintances who are not co-authors. Yellow circles represent each author's ten closest authors in academic cosine similarity. The lengthier edges represent longer distances.

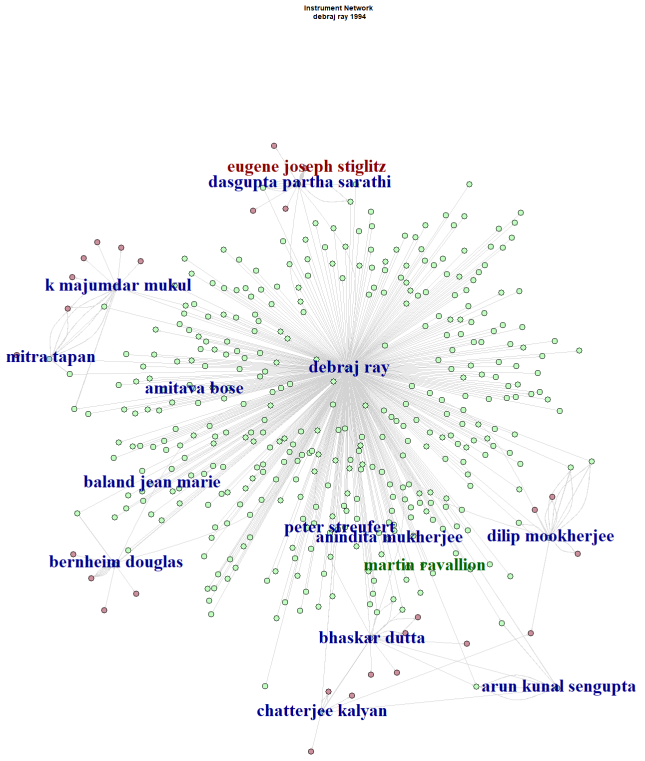


(b) Distribution of academic cosine similarity between Ariel Rubinstein and all other economists. A subset of his co-authors are marked along the x axis by their names; the density of his non-acquaintances appears in pink; the densities of his acquaintance sets appear in blue ($n = 5$), red ($n = 10$), and green ($n = 20$). The vertical dashed line represents the location of Rubinstein's tenth most similar author.

Figure 5: Illustration: Ariel Rubinstein's and Martin Osborne's local peer network, and distribution of Ariel Rubinstein's academic similarities.

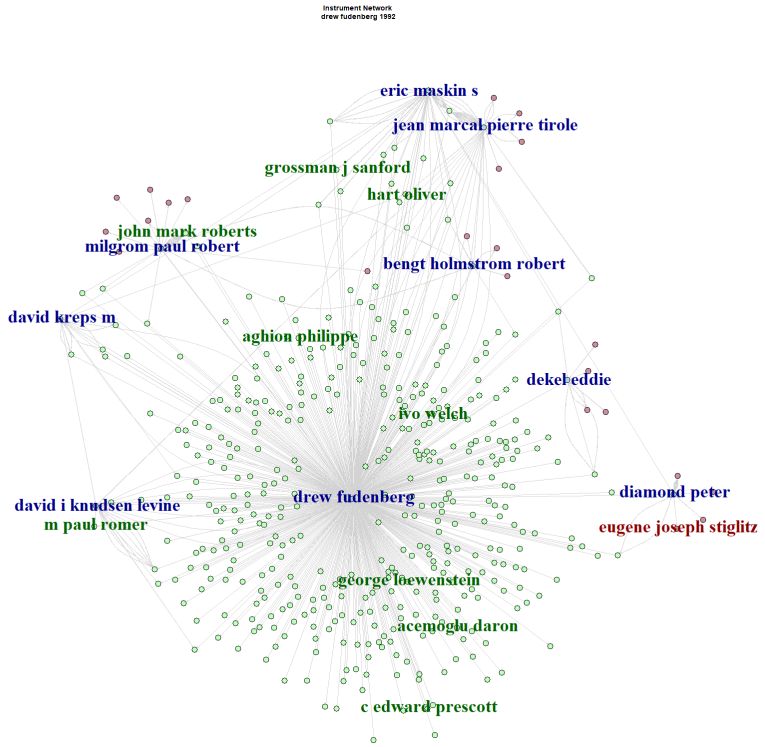


(a) Debraj Ray's network, 1993

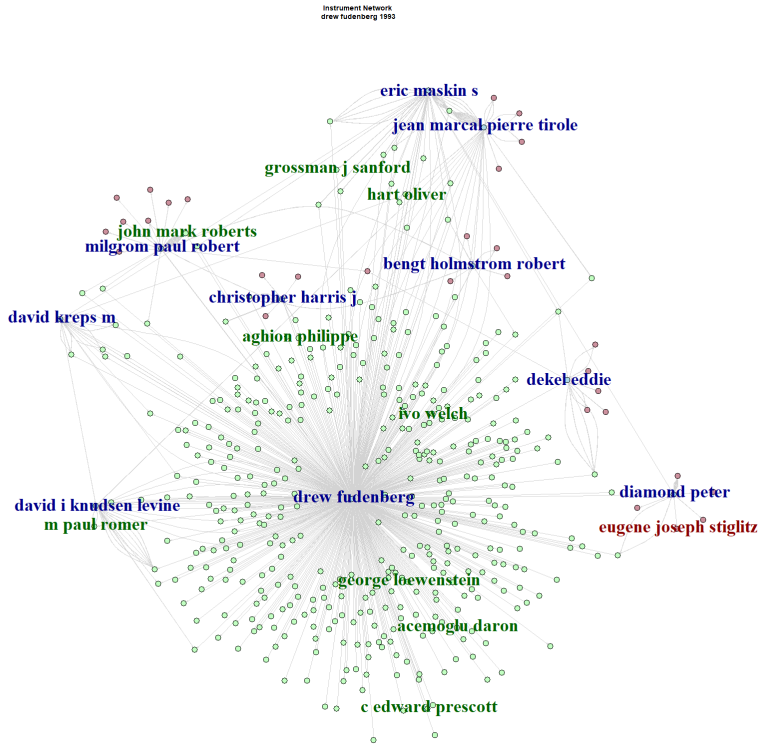


(b) Debraj Ray's network, 1994

Figure 6: Example illustrating the instrumental variables variation induced by co-authors of co-authors who are not acquaintances of an author. Co-authors appear in blue, acquaintances appear in green, and non-acquaintances appear in pink.



(a) Drew Fudenberg's network 1992



(b) Drew Fudenberg's network 1993

Figure 7: Example illustrating the instrumental variables variation induced by co-authors of co-authors who are not acquaintances of an author. Co-authors appear in blue, acquaintances appear in green, and non-acquaintances appear in pink.

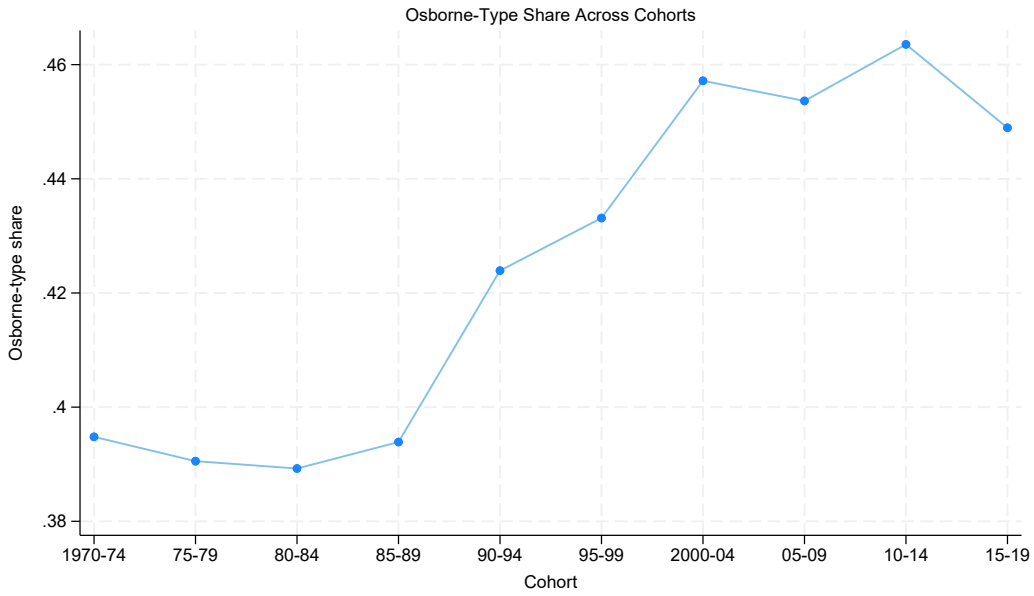


Figure 8: Osborne Type Share across Cohorts. Share of authors assigned to Osborne's community, by 5-year cohorts of economists based on the community detection estimates based on the ten-closest acquaintance set definition.

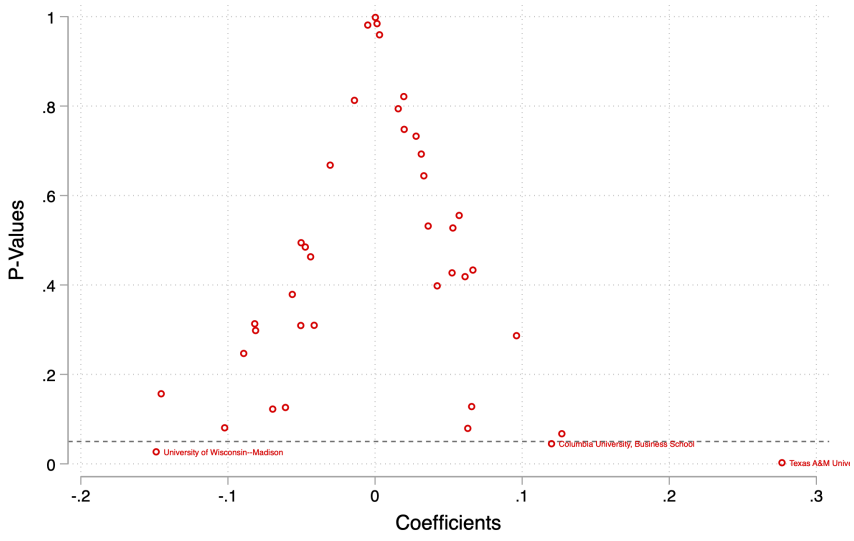
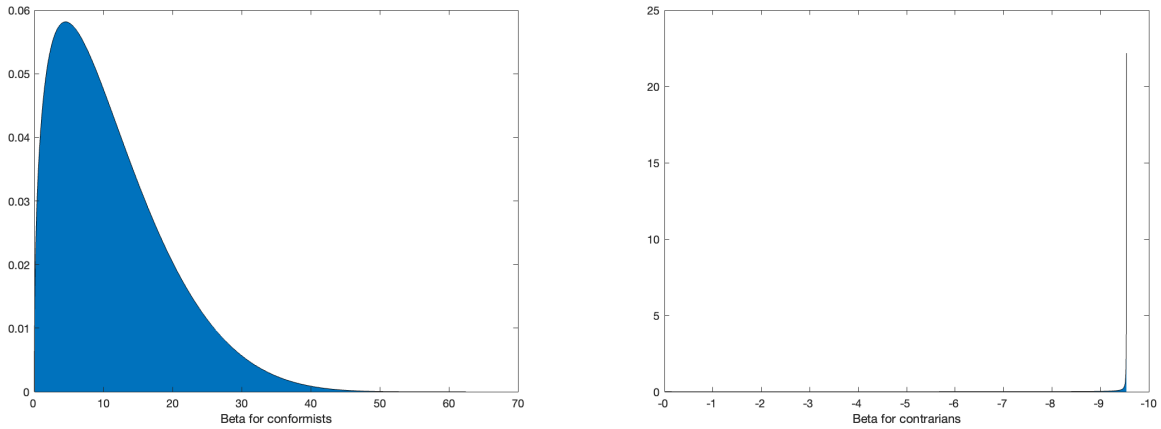


Figure 9: University affiliations and the Osborne-type dummy. Distribution of coefficient sizes and p-values by university to predict the Osborne-type dummy in a regression of 1,868 unique authors in 39 academic departments and 2,592 authors-x-department of the form: Osborne type dummy$_i = a + \beta$Univeristy j dummy$_i + \epsilon_i$. The dashed line represents a p-value of 0.05.



(a) $a = 1.4$, $b = 6.3$, $S = 62.4$.

(b) $a = 0.75$, $b = 0.02$, $S = 9.5$.

Figure 10: Peer effects. Estimated Beta distributions of peer effect heterogeneity for conformist and contrarian economists.

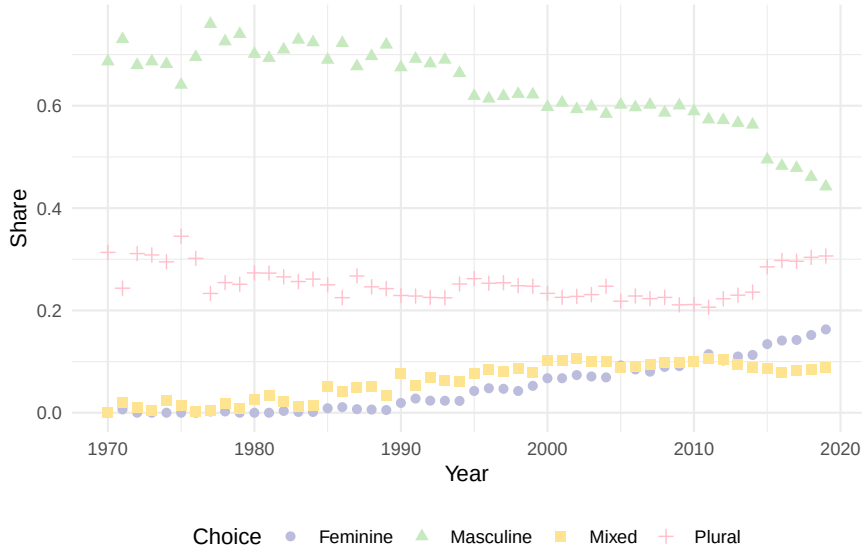


Figure 11: Time-series model fit. The figure plots the aggregate distribution of pronoun choices over time from simulated choices based on the estimated parameters from Table 9. As starting values for the peer influences, the simulation uses the observed average 1970-1974 distribution of choices.

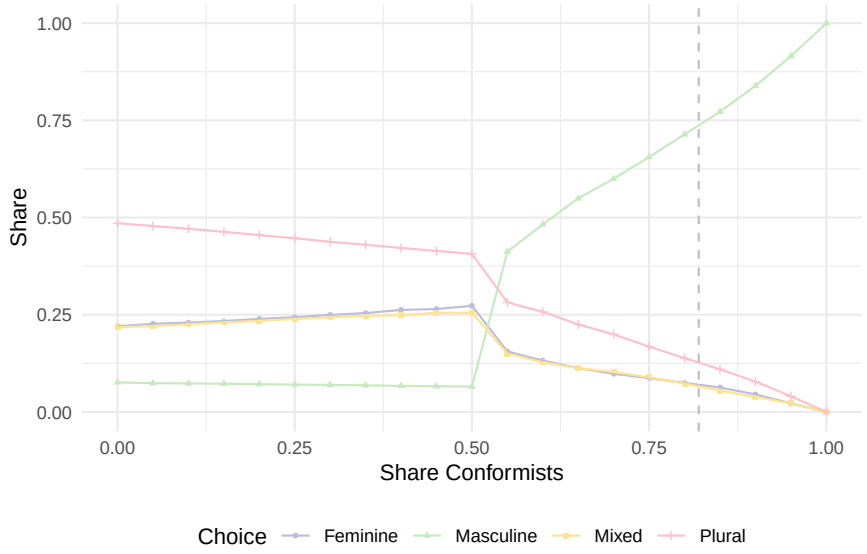


Figure 12: Counterfactual simulations under alternative distributions of the psychological types. The figure plots the aggregate distribution of end line (2014-2019) pronoun choices from simulated choices based on the estimated parameters from Table 9, under varying shares of conformists in the population. As starting values for the peer influences, the simulation uses the observed average 1970-1974 distribution of choices.

10 Online Appendix I: Additional Tables and Figures

From/To	Masculine (1)	Feminine (2)	Plural (3)	Mixed (4)
Single authored				
Masculine	0.63	0.03	0.21	0.13
Feminine	0.14	0.43	0.20	0.23
Plural	0.32	0.07	0.50	0.11
Mixed	0.28	0.11	0.19	0.42
70s cohort				
Masculine	0.65	0.02	0.22	0.10
Feminine	0.22	0.27	0.21	0.30
Plural	0.41	0.04	0.44	0.12
Mixed	0.33	0.07	0.24	0.35
80s cohort				
Masculine	0.54	0.04	0.25	0.17
Feminine	0.19	0.27	0.27	0.27
Plural	0.31	0.07	0.49	0.13
Mixed	0.31	0.11	0.22	0.36
90s cohort				
Masculine	0.48	0.07	0.24	0.20
Feminine	0.19	0.27	0.25	0.29
Plural	0.25	0.10	0.50	0.15
Mixed	0.28	0.14	0.20	0.38
00s cohort				
Masculine	0.46	0.09	0.23	0.22
Feminine	0.17	0.33	0.24	0.25
Plural	0.23	0.13	0.50	0.15
Mixed	0.26	0.17	0.18	0.39
10s cohort				
Masculine	0.41	0.12	0.24	0.23
Feminine	0.16	0.38	0.22	0.24
Plural	0.21	0.13	0.52	0.14
Mixed	0.24	0.19	0.18	0.39

Table A.11: Sub-group transition matrices. Transition matrices for single-authored to single-authored papers, and for different cohorts of authors.

	Osborne dummy
Woman	-0.021
	(0.053)
Woman × 1980	-0.050
	(0.061)
Woman × 1990	-0.009
	(0.057)
Woman × 2000	-0.001
	(0.055)
Woman × 2010	0.003
	(0.055)
Obs.	29302

Table A.12: The table presents the coefficients and standard errors from a cross-sectional linear regression at the author level, of the Osborne dummy on a dummy for whether the author is a woman, and interactions of it with cohort dummies.

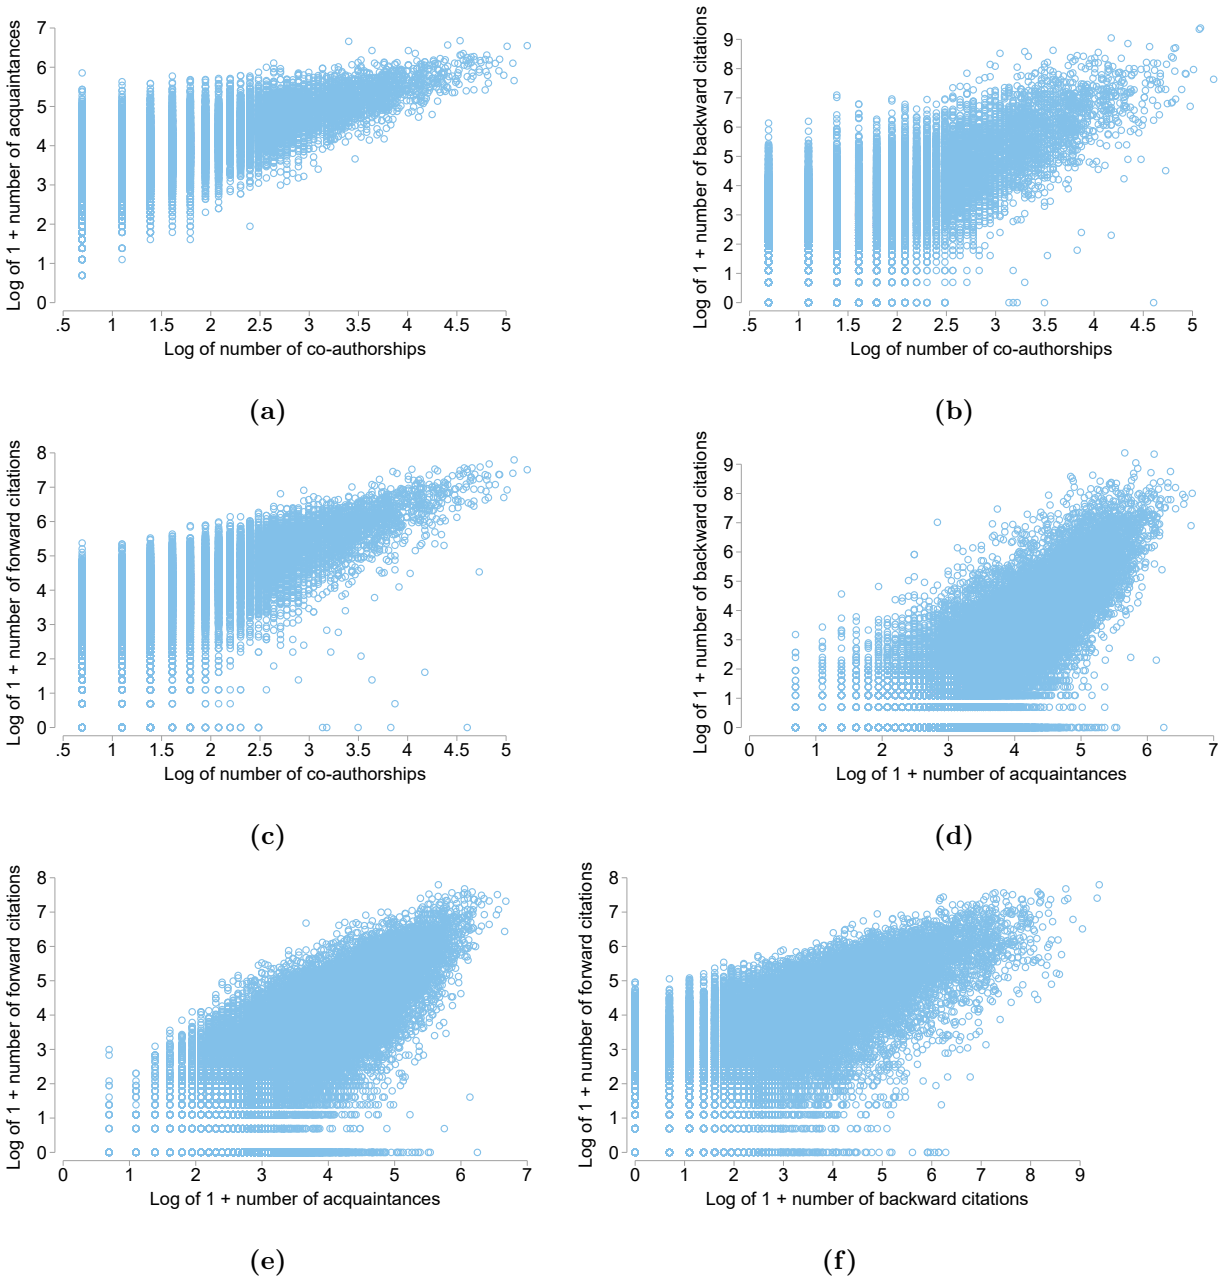
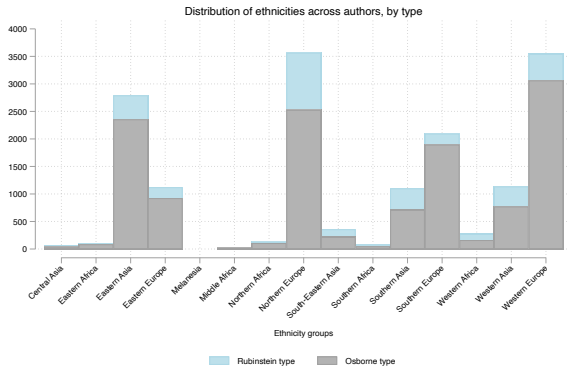
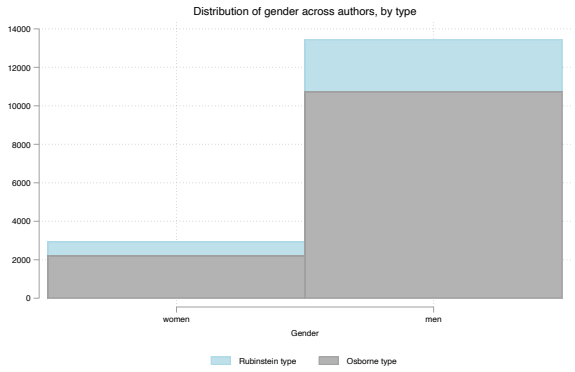


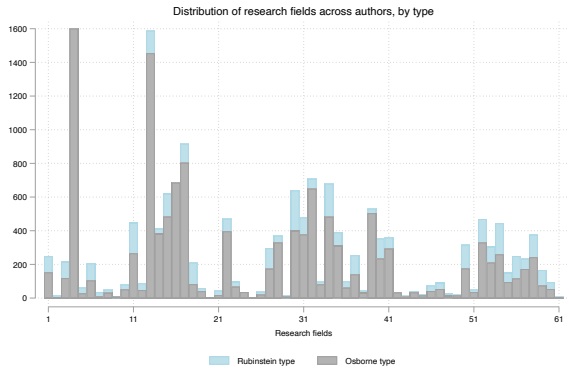
Figure A.13: Degree distributions across networks.



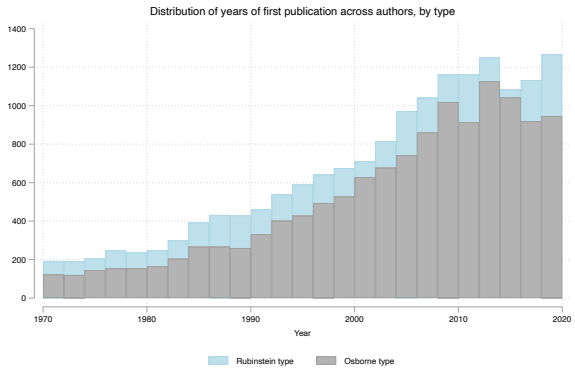
(a) Ethnicity



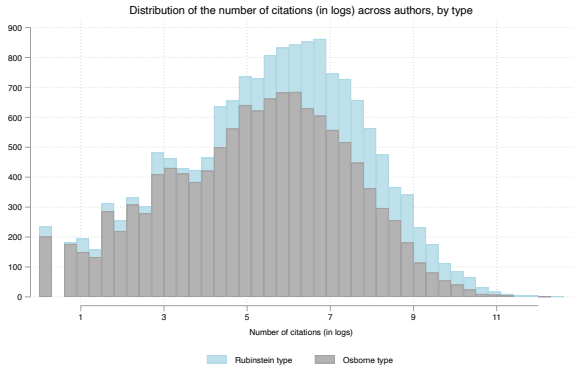
(b) Gender



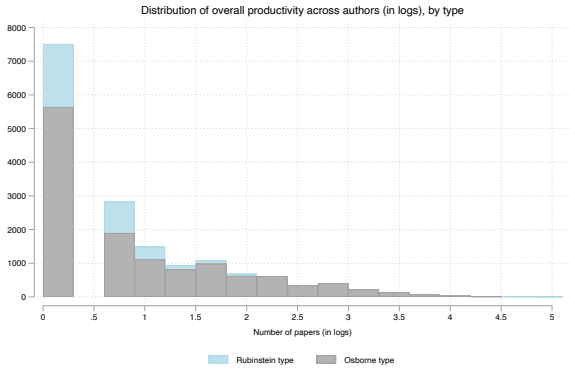
(c) Fields



(d) First publication



(e) Citations



(f) Productivity

Figure A.14: Distributions of author characteristics by community type assigned.

11 Online Appendix II: Methodological Details

11.1 Selection of the sample of articles and authors

We use several sources to put together the sets of articles and authors that underlie our study. From *Jstor* and *Crossref* we obtained the metadata and the full texts of a large set of papers from Economics and Economics-related academic journals. We obtained the *Jstor* data under a data user agreement for the project. We obtained the *Crossref* data using their API.[32]

This resulted in 710 thousand articles. The set is over-inclusive, however. It contains papers in all fields of Economics, whereas our purpose is to put together a set of economic theory articles only. We implement a layered procedure to filter out articles unlikely to be theoretical, and to make sure we keep articles likely to be theoretical.

1. We exclude articles with corrupted metadata:
   - Missing a title.
   - Missing authors.
   - Missing the articles' text. These are articles for which our *Crossref* API retrieval generated a line of metadata but no associated article text. We inspected the list titles of this set of articles, and found 849 that we clearly identified as economic theory papers. We proceeded to directly retrieve the text of these articles, and included them back.
2. We exclude any article whose metadata suggests it is not a standard academic paper. This includes a reference to any of the following labels:

```
"Note from the editor"                  "Photograph"
"Meeting of the econometric society"    "Meetings of the econometric society"
"Accepted Manuscripts"                  "List of members"
"Announcement"                          "Announcements"
"Award"                                 "Awards"
"Front matter"                          "Back Matter"
"Book review"                           "Book reviews"
"Call for papers"                       "Distinguished fellow"
"Referees"                              "Editorial"
"Editor"                                "Election of fellows"
"Errata"                                "Erratum"
"Addendum"                              "Correction:"
"Correction to:"                        "Retracted Article"
"Corrigendum"                           "European meeting"
"Fellows"                               "Foreward"
```

[32] See https://www.crossref.org/education/retrieve-metadata/rest-api/. We used the R package `crminer` to retrieve the data. This package is no longer maintained, and to our knowledge, Crossref discontinued its open-access full-text retrieval service as of December 2020 -after we accessed it-.

```
"In memoriam"                          "Obituary"
"Report of the committee"              "Report on the adhoc committee"
"Report of the director"               "Report of the editor"
"Report of the managing editor"        "Report of the representative"
"Report of the secretary"              "Report of the treasurer"
"Submission"                           "Report of the President"
"Thesis titles"                        "Author index"
"Discussion"                           "Preface"
"Foreword"                             "Index"
"Comment"                              "Contributors"
"Abstracts"                            "Noticeboard"
"IMACS"                                "Reply"
"Note"                                 "Rejoinder"
"Presidential address"                 "Hardback"
"Hardcover"                            "Paperback"
"Actuarial Vacancy"                    "Secretary-Treasurer"
"Secretary/Treasurer"                  "Treasurer"
"ISBN"                                 "pp\\."
"Conference"                           "Symposium"
"Verlag"                               "pages"
"Tribute"                              "(Eds)"
"Listing Service"                      "Content of Volume"
"Contents of Volume"
```

3. We exclude all articles from academic journals that are either exclusively econometric or statistical, or from unrelated fields. Below is the list of journals whose articles we exclude:

```
"Econometric Theory"
"Econometrics Journal"
"Journal of Applied Econometrics"
"Journal of Econometrics"
"Physica A: Statistical Mechanics and its Applications"
"Statistics & Probability Letters"
"Stochastic Processes and their Applications"
"Applied Energy"
"Energy"
"Resources and Energy"
"Renewable Energy"
"The Electricity Journal"
"Marine Policy"
"Computational Statistics & Data Analysis"
"Mitigation and Adaptation Strategies for Global Change"
"Journal of Classification"
"World Patent Information"
```

"Social Indicators Research: An International and Interdisciplinary Journal for Quality-of-Life Measurement"
"Journal of Multivariate Analysis"
"Metrika: International Journal for Theoretical and Applied Statistics"
"Statistical Papers"
"Annals of the Institute of Statistical Mathematics"
"Journal of the Royal Statistical Society Series A"
"Journal of the Royal Statistical Society. Series C (Applied Statistics)"
"Journal of Time Series Analysis"
"Statistical Methods & Applications"
"Applied Mathematics and Computation"
"Mathematics and Computers in Simulation (MATCOM)"
"Global Finance Journal"
"Children and Youth Services Review"
"European Journal of Operational Research"
"Mathematical Methods of Operations Research"
"Mathematics of Operations Research"

4. We directly included in our final set all articles from strictly economic theory journals:

   "Journal of Economic Theory"
   "American Economic Journal: Microeconomics"
   "Economic Theory"
   "Games and Economic Behavior"
   "International Journal of Game Theory"
   "Games"
   "Journal of Public Economic Theory"

5. For all other articles which had not been filtered out at this stage, we implement an algorithm to classify them as likely theoretical. For this purpose, we constructed a list of microeconomics keywords and a list of econometrics keywords.

   The list of microeconomics keywords is:

   game, player, utility, coalition, equilibrium,
   equilibria, rational, preference, core, Bayesian,
   pricing, welfare, marginal cost, theoretic, induction,
   signalling, strategic, bargaining, proposal, dynamic,
   Markov, subgame, monopoly, duopoly, oligopoly, cooperation,
   free rid, punish, design, contract, first best, second best,
   model, theory, theories, theoretical, auction, bid,
   dominance, risk, payoff dominant, backward induction, Cournot,
   Stackelberg, Nash, Aumann, unique, existence, multiplicity,
   pure, mixed, coordination, hawk, dove, battle of the sexes,
   battle of the sex, matching pennies, prisoner, efficient,

`efficiency, evolutionary, replicator, dynamics, stable,`
`opponent, ambiguity aversion, strategies, payoffs,`
`expected utility, common knowledge, match, beliefs, intuitive`
`criterion, fixed point, delay, market design, zero-sum,`
`n-person, linear programming, Marshallian, compensated variation,`
`transitive, transitivity, club, Rules of thumb, rule of thumb,`
`Shapley value, Axiom, Axiomatic, Normal form, Extensive form,`
`Information set, Impossibility, Information structure,`
`private information, asymmetric information, moral hazard,`
`adverse selection, surplus, incentive constraint,`
`participation constraint, transferable utility, quasi-linear`

The list of econometrics keywords is:

`estimator, instrument, asymptotic variance, regression,`
`two stage least square, maximum likelihood,`
`generalized method of moments, multiple test, delta method,`
`continuous mapping theorem, measurement error, moment condition`

(a) We include any paper containing at least 250 microeconomics keywords and no econometrics keywords.

(b) We include any paper satisfying all of the following criteria:

- Contains the word *proof* in its text.
- Contains at least ten microeconomics keywords.
- Contains ten times more microeconomics keywords as econometrics keywords.

(c) We then identify all authors from papers from (a) and (b), and among the remaining not-yet-included papers, we include those which satisfy the both of the following conditions:

- It includes authors from this list.
- it has ten times more microeconomics theory keywords as econometrics specific keywords, or has zero econometrics specific keywords.

This concludes the first component of the selection of papers into our sample, and yields 70062 articles written by 48626 authors.

6. At this stage, some of these 48626 author names correspond to differing spellings of the name of the same underlying author. We implemented an algorithm to find the alternative spellings of the same author, to then collapse these alternative spellings into a single author. First, we compute the frequencies of each name component (e.g., a first name, a last name, etc.) among all author names. We also extract the initials of each full name. We then identify, for each author, his least common name component (we call it the rare component). For example, for *Jean Marcal Tirole*, *Marcal* is its rare component, as its frequency is the smallest among the three components of this name. Next we split the sample of author full names into two sets. A set A of authors whose

rare component is unique in the data set, and none other of the components of their names are a rare component of any other author, and a set B with its complement.

The uniqueness of at least one word in the names of authors in set A implies they are highly unlikely to have duplicates. Set A has 8137 authors. In contrast, authors in set B have a rare component that is not unique in the data set. For each author $i \in B$, we produce a list of potential duplicates $D(i) = \{j, k, ...\} \subset A \cup B$ containing the author identifiers of each author sharing i's rare component. We then compare the initials of i's name to the initials of the names of every element of this potential match list to thin out these lists as follows: if j's initials are not a subset, a super set, or identical to the initials of i, we exclude j from the list. If the resulting match list for i is empty, we consider i to have no duplicates and hence to be unique. We identify 18413 authors as unique in this step.

For authors i for whom this procedure yields non-empty potential match sets $D(i)$, we further make pairwise comparisons of each of the name components of i to each of the name components of j with overlapping initials. If there is not at least one identical pairing among all these comparisons, we exclude j from the list in an additional thinning step. If the resulting match list for i is empty, we consider i to have no duplicates and hence to be unique. We identify 6336 authors as unique in this step. This leaves us with $15740 = 48626 - 8137 - 18413 - 6336$ author names i with potential duplicates $D(i)$, with corresponding initials and at least one identical name component from set B.

We then move to compare them to their potential duplicates using information about their articles. To do this, we first take the titles of the articles of each author i, and retrieve *ChatGPT* embeddings for each title separately, $\mathbf{e}_{ia}$, and for the grouping of all the titles of the author's articles, $\tilde{\mathbf{e}}_i$. For each pair of potential duplicate authors we compute the cosine similarity between each pairing of their articles and find the highest of these cosine similarities, s_{ij}^{max}. For each pair of potential duplicates authors we compute the cosine similarity between their grouped-titles embedding, $\tilde{s}_{ij}$. We then apply the following rule:

(a) If authors i and j share the same rare component (stronger signal), and $min\{s_{ij}^{max}, \tilde{s}_{ij}\} \geq 0.8$, consider i and j to be the same author.

(b) If authors i and j share a name component that is not the rare one for one of the authors (weaker signal), and $min\{s_{ij}^{max}, \tilde{s}_{ij}\} \geq 0.9$, consider i and j to be the same author.

(c) Otherwise, consider i and j to be unique distinct authors.

We chose the cutoff values for these rules by inspecting the sample and trade-off type 1 and type 2 errors as best as possible. In this way, we incorporate information from both the pair of authors' names and from the similarity in their articles, to assess whether they are actually the same individual.

For the remaining set of names i and potential duplicates $D(i)$ we find the most similar duplicate of i, $m_i = argmax_{j \in D(i)} \tilde{s}_{ij}$. We then find the most similar author to m_i: m_{m_i}.

If $m_i \neq m_{m_i}$, i.e., if the most similar duplicate of i does not have i as its most similar duplicate too, we consider them to be distinct authors unless $\tilde{s}_{im(i)} > 0.85$. Otherwise, we classify them as the same author. This final step is particularly useful for a handful of cases with a multiplicity of differing but closely similar name variations.

For the top 200 authors in our data set based on citations, we manually checked for alternative spellings of their names, and collapsed the duplicates accordingly. At this point we are left with 46655 unique authors.

7. Finally, we exclude articles missing their publication date, or with publication dates prior to 1970 or posterior to 2019. We also exclude articles that do not use any third person pronouns as described in subsection 11.2, and articles that do not have at least one known author matched to it.

This concludes our construction of the sample of articles and authors, and yields 73099 articles and 38046 unique authors.

11.2 Classification of the pronoun use style of articles: Allen NLP correferencing

Our methodology demands that we classify the writing style of each article as it relates to the gender choices for its third person pronouns. We rely on the *Allen* natural language processing (NLP) package, a state-of-the-art neural network model.[33] For each paper, we identify every instance of one of the following third-person pronouns: Masculine:

```
he, him, his, himself.
```

Feminine:

```
she, her, hers, herself.
```

Plural:

```
they, them, their, theirs, themselves.
```

Mix:

```
he or she, him or her, his or her, himself or herself, he and she,
him and her, his and her, himself and herself.
```

For each identified pronoun, we extract the sentence containing the pronoun, and the sentences preceding and following it. We then run the *Allen NLP* correferencing model on this text segment. This model relates the pronoun to its corresponding noun within the segment. For example, if we feed it the sentence "John ate an apple and he liked it", *Allen NLP* will indicate that "he" refers to John, and that "it" refers to apple. Figure A.15 illustrates the form of the *Allen NLP* output, in a paragraph from an article in our sample.

[33] See https://demo.allennlp.org/coreference-resolution.

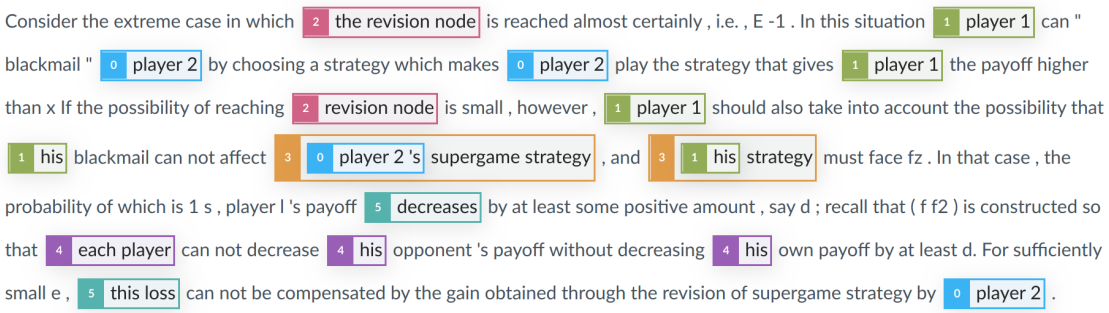


Figure A.15: *Allen NLP* correferencing example.

After parsing every segment involving a pronoun in an article, we obtain a list of proper nouns and their corresponding correferenced third-person pronouns in the article. *Allen NLP* is known to achieve at least a 75 percent accuracy in standard English text. At the paper level, our manual checks suggest an error rate of almost zero.

In a next step, we use a list of keyword economic agent nouns, to select the *Allen NLP* correferenced nouns in each paper that correspond to economic agents the articles are referring to. We use the following list:

```
'individual', 'worker', 'agent', 'principal', 'loser',
'representative','[pl]ayer', 'trader', 'competitor', 'winner',
'citizen', 'messenger', 'manufacturer', 'investor', 'bank',
'government', 'criminal', 'member', 'researcher', 'opponent',
'group', 'respondent', 'party', 'incumbent', 'buyer', 'legislator',
'officer', 'prisoner', 'insured', 'insurance', 'owner', 'lender',
'challenger', 'cooperator', 'employer', 'customer', 'participant',
'borrower', 'mover', 'recipient', 'household', 'innovator', 'leader',
'rival', 'follower', 'contestant', 'intermediaries', 'voter',
'dictator', 'ceo', 'monopolist', 'migrant', 'candidate', 'manager',
'peer', 'user', 'trustee', 'oligopolist', 'employee', 'firm',
'regulator', 'person', 'maker', 'auctioneer', 'type', 'intruder',
'outsider', 'insider', 'people', 'dealer', 'entrepreneur',
'policymaker', 'nature', 'negotiator', 'neighbo[r]', 'executive',
'physician', 'generation', 'child', 'parent', 'newcomer', 'friend',
'professional', 'retailer', 'resident', 'student', 'subject',
'seller', 'partner', 'bidder', '[c]onsumer', 'organization',
'those who', 'sender', 'receiver', 'stockholder', 'team', 'speculator',
'supplier', 'producer', 'labourer', 'laborer', 'landholder', 'farmer',
'developer', 'creditor', 'politician', 'planner', 'arbitrageur',
'committee', 'board', 'bargainer', 'herder', 'defendant', 'plaintif',
'jury', 'jurist', 'juror', 'judge', 'colleague', 'faculty', 'scientist',
'analyst', 'applicant', 'baron', 'bureaucrat', 'contractor',
'decision - maker', 'decisionmaker', 'decisions makers', 'entrant',
'expert', 'landlord', 'merchant', 'mutant', 'offender', 'peasant',
'proposer', 'purchaser', 'responder', 'teacher', 'venture capitalist',
```

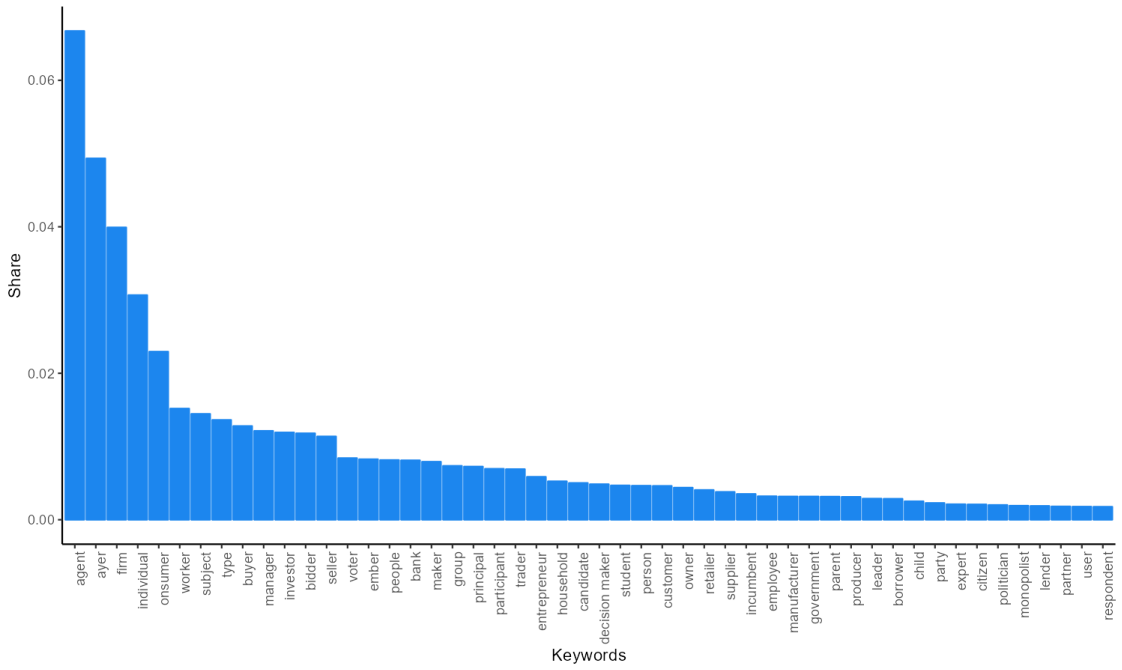


Figure A.16: Distribution of agent nouns used for co-referencing across articles: top 50.

`'tortfeasor', 'commuter', 'insurer'.`

After identifying all instances of pronoun use referring to any of the agent nouns listed above, we count the number of times masculine, feminine, plural, or a combination, are used in each paper to refer to them. Figure A.16 presents the distribution of these agent nouns across the full sample of article texts, for the top 50 most frequently used agent nouns.

We classify an article as masculine if it only uses masculine pronouns. We classify an article as feminine if it only uses feminine pronouns. We classify an article as plural if it only uses plural pronouns. We classify an article as mixed if it uses mixed pronouns, or a combination of more than one type of pronoun.

11.3 Measurement of the relative spatial location of authors: *Author2vec*

To identify a set of plausible coauthors for each author in our sample, we adapted the *Word2vec* algorithm to our setting. *Word2vec* is a widely used algorithm in computer science designed to capture semantic relationships between words based on their co-occurrence patterns in a body of text (corpus). It is based on the distributional hypothesis proposition in linguistics, according to which words appearing in similar contexts tend to have similar meanings. Within a given corpus (e.g., the congressional record), it uses the relative frequencies with which pairs of words appear near each other (right before or after, within a few words of each other, etc.) to assign a high-dimensional vector of real numbers to each

word –referred to as the word's *embedding*–.[34] We denote word i's embedding by $\mathbf{e}_i$. An embedding contains cardinal information about the word's meaning in relation to all other words in the corpus: words that are closer to each other in this vector space, say using a Euclidean distance norm, are deemed to be closer to each other in meaning, because the relative frequencies with which they appear near other words is similar.

Consider a word w_j in some sentence, and refer to it as the center word. Consider other words in the same sentence found at most m[35] words away from w_j, and refer to them as context words. Denote this set as $M(j)$. *Word2vec* allows for each word j to have an embedding as center word, $\mathbf{e}_j^c$, and an embedding as context word $\mathbf{e}_j^o$. *Word2vec* defines the conditional probability of observing context word w_k given center word w_j using the *softmax* function as

$$\mathbb{P}(w_k|w_j) = \frac{\exp(\mathbf{e}_k^{o\prime}\mathbf{e}_j^c)}{\sum_\ell \exp(\mathbf{e}_\ell^{o\prime}\mathbf{e}_j^c)}$$

Making the dot product between context word w_k and center word w_j large relative to all other words in the corpus makes this probability high.

Word2vec chooses the collection of vectors $\{\mathbf{e}_j^o, \mathbf{e}_j^c\}_{j=1}^W$ for all words in the corpus that maximizes the joint likelihood of observing the actual context-center pairs:

$$L(\theta) = \prod_{j=1}^{W} \prod_{k \in M(j)} \mathbb{P}(w_k|w_j)$$

The solution to this problem minimizes the difference between the predicted conditional probabilities and the actual distribution of word pairings in the corpus. In a final step one can average the estimated center and context embeddings of each word to obtain a single embedding for the word.

Word2vec is, implicitly, a network-based model where words are nodes, and edges between words exist when two words are near each other in the corpus –how near being a parameter chosen by the researcher–. The idea we propose here is to rely on the same logic, applied to the social network of economists in our sample, to measure 'academic similarity' across authors. We call this algorithm *Author2vec*. Authors play the role of words, cross-citation relationships play the role of edges between them, and we compute an embedding vector for each author[36]. Two authors with close embeddings will be authors who cite and are cited by similar subsets of other authors, in the same way that words with close embeddings are words that appear near similar subsets of other words. In this sense, such authors are nearby in 'academic' space, and we will rely on this academic distance to restrict the set of authors that could feasibly be co-authors of a given author.

To implement our *Author2vec* methodology we transform each article a in our data set into a vector $\mathbf{v}_a$ of author identifiers that includes identifiers for all authors that either co-authored the paper or that are cited in the paper. Each such vector is analogous to a sentence in standard *Word2vec*. The collection of all such vectors $\{\mathbf{v}_1, \mathbf{v}_2, ..., \mathbf{v}_{N_a}\}$ constitutes

[34] Large language models such as *ChatGPT*, for example, rely on a corpus that may include all of the internet, and on embeddings of many thousands of dimensions.

[35] m is a radius chosen by the researcher. If $m = 1$, for example, we only consider the word directly preceding and the word directly succeeding w_j as context words.

[36] In practice we allow for 100-dimensional embeddings for the authors.

our corpus. We define a pair of authors to be 'near' if they appear in the same article vector. We can then use the frequencies with which each author is 'near' every other author within our corpus of articles in exact analogy to how *Word2vec* uses the frequency with which a given word appears before or after (near) every other word within the corpus of text.

We rely on the *Microsoft Academic Graph* (MAG)[37] and *Jstor* data sets to retrieve network-related information about the set of authors in our sample, including co-authorship relationships and forward and backward citation relationships.

11.4 Construction of the acquaintance sets

We rely on the author embeddings from our *Author2vec* methodology to compute the cosine similarity (dot product of two vectors divided by the product of their lengths) between each pair of authors in our sample, $s_{i,j}$, as a scalar measure of academic proximity[38]:

$$s_{i,j} = \frac{\mathbf{e}_i' \mathbf{e}_j}{|\mathbf{e}_i||\mathbf{e}_i|}$$

Our premise is that pairs of authors far from each other in this academic space are effectively unable to consider each other as potential co-authors. We compute an acquaintance set of potential co-authors for each author, $Q_n(i)$, as follows: we take the union of the n closest authors to author i, all co-authors of author i, and the n closest authors to each of i's co-authors. We then exclude from this set any author who does not overlap in his productive years –defined as the range of years between three years before the author's first publication and five years after the author's last publication–, with author i. By construction, $Q_n(i)$ includes all authors who did co-author with i at some point and a number of other authors who did not, but who are close enough in academic space that it is likely i could have considered them as co-authors. Our benchmark estimates use acquaintance sets with $n = 10$, but we also set $n = 5$ or 20 in alternative specifications.

11.5 Measurement of covariates

11.5.1 Assignment of sub-fields for authors: *ChatGPT* embeddings

Co-authorship decisions are likely influenced, among other characteristics, by the overlap in the sub-fields of study of authors. We assign sub-fields of specialization to the authors in our sample as follows: first, we borrow the *Journal of Economic Literature* (JEL) fields classification, and select a subset of the JEL fields which we deem relevant in our context. The following is the list of JEL fields we use:

- C6 Mathematical Methods • Programming Models • Mathematical and Simulation Modeling
- C7 Game Theory and Bargaining Theory

[37] See https://www.microsoft.com/en-us/research/project/microsoft-academic-graph.

[38] Cosine similarity is the most commonly used distance measure in the network science-large language models literature.

- C9 Design of Experiments
- D1 Household Behavior and Family Economics
- D2 Production and Organizations
- D3 Distribution
- D4 Market Structure, Pricing, and Design
- D5 General Equilibrium and Disequilibrium
- D6 Welfare Economics
- D7 Analysis of Collective Decision-Making
- D8 Information, Knowledge, and Uncertainty
- D9 Micro-Based Behavioral Economics
- E2 Consumption, Saving, Production, Investment, Labor Markets, and Informal Economy
- E3 Prices, Business Fluctuations, and Cycles
- E4 Money and Interest Rates
- E5 Monetary Policy, Central Banking, and the Supply of Money and Credit
- E6 Macroeconomic Policy, Macroeconomic Aspects of Public Finance, and General Outlook
- E7 Macro-Based Behavioral Economics
- F1 Trade
- F3 International Finance
- G1 General Financial Markets
- G2 Financial Institutions and Services
- G3 Corporate Finance and Governance
- G4 Behavioral Finance
- G5 Household Finance
- H1 Structure and Scope of Government
- H2 Taxation, Subsidies, and Revenue
- H3 Fiscal Policies and Behavior of Economic Agents

- H4 Publicly Provided Goods
- H5 National Government Expenditures and Related Policies
- H6 National Budget, Deficit, and Debt
- H7 State and Local Government • Intergovernmental Relations
- H8 Miscellaneous Issues
- I1 Health
- I2 Education and Research Institutions
- I3 Welfare, Well-Being, and Poverty
- J. Labor and Demographic Economics
- K. Law and Economics
- L1 Market Structure, Firm Strategy, and Market Performance
- O1 Economic Development
- O2 Development Planning and Policy
- O3 Innovation • Research and Development • Technological Change • Intellectual Property Rights
- O4 Economic Growth and Aggregate Productivity
- P. Political Economy and Comparative Economic Systems
- R. Urban, Rural, Regional, Real Estate, and Transportation Economics
- Z1 Cultural Economics • Economic Sociology • Economic Anthropology

We then retrieve the *ChatGPT* embedding corresponding to all the words in the description of each of these fields, including the text descibing its subfields.[39] This gives us an embedding for each field j, $\bar{\mathbf{f}}_j$, with $j = 1, ..., J$. In parallel, for each author i in our sample we create a collection K_i of the words in the titles of all of i's articles, and the words in the titles of all papers cited in i's articles. Next we retrieve the *ChatGPT* embedding for the

[39] We retrieve ChatGPT-3 embeddings of 1536 dimensions, based on their text-embedding-ada-002 model. See https://openai.com/blog/new-and-improved-embedding-model. Because ChatGPT's embeddings are estimated for a large corpus of English text, they are ideal as measures of relative similarity between common-use words. One of the main advantages of LLM word embeddings is their cardinal nature, allowing arithmetic operations that preserve relative meanings. As an example often used in this literature, subtracting the embedding for the word *man* from the embedding for the word *king*, and then adding the embedding for the word *woman* yields an embedding that is remarkably close to the embedding for the word *queen*.

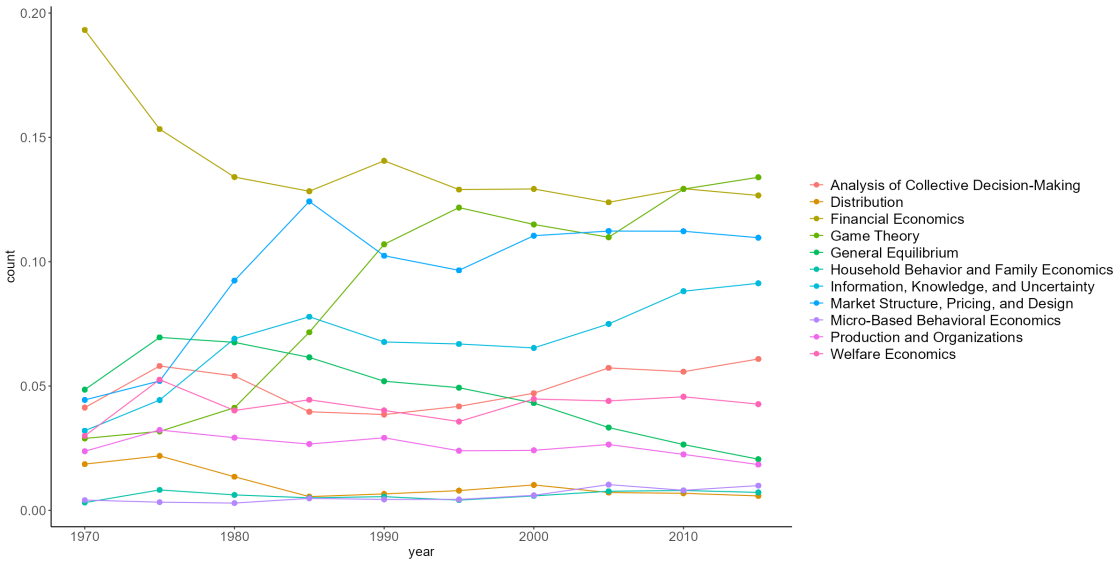


Figure A.17: Distribution sub-fields by 10-year cohorts.

collection of all words in K_i. This gives us an average embedding for author i, $\overline{\mathbf{g}}_i$. Next we compute cosine similarity distances between each author and each field,

$$\sigma_{i,j} = \frac{\mathbf{g}_i' \mathbf{f}_j}{|\mathbf{g}_i||\mathbf{f}_j|}.$$

Finally, we assign to each author the three sub-fields with the smallest cosine similarity distances and use those to create dummy variables indicating sub-fielf membership.

In Figure A.17 we plot the distribution of sub-fields by 5-year cohorts of articles. Most fields have remained stable, with some exceptions: "Financial Economics" which fell from 20 percent to 13 percent in the 1970s and has remained stable since, and General Equilibrium has fallen steadily from 7 percent to 3 percent. In contrast, Game Theory and Market Structure, Pricing and Design grew rapidly in the 70 and 80s from less than 5 percent to around 12 percent today each. Information, Knowledge, and Uncertainty has also grown from 3 to 9 percent.

11.5.2 Classification of the of ethnic origin of authors: *Namsor*

We rely on the authors' full names we obtained directly from the articles in our data set to assign an ethnic origin to each author. We do this using *Namsor*[40], a software tool specialized in identifying the likely regions of origin of proper names and last names from cultures all around the world. For each component of an author's name –first name, middle name, last name– Namsor reports a most likely origin at the sub-region level (e.g., Western Europe, South-east Asia, Middle East, etc.). As the ethnic origin of author i, we assign the modal sub-region reported by Namsor across all of the author's name components. For the small subset of cases with ties, we relied on *ChatGPT* prompts containing Namsor's guesses, and

[40]See https://namsor.app.

retrieved *ChatGPT* best guess response. .

11.5.3 Classification of the gender of authors: Genderize package in R

We rely on the authors' first names we obtained directly from the articles in our data set to assign a gender to each author. We do this using the *Genderize* package in R,[41] a software tool that has been trained on a large corpus of text as a probabilistic gender classifier for first names. We face one challenge: first and last names appear in no particular order. Sometimes first names appear before last names, and sometimes the other way around. Thus, we proceeded by genderizing each component of an author's full name. For example, we asked the package to assign a gender to both "Debraj" and "Ray" separately. We then classified the authors as follows: if both components were assigned the same gender, we assigned that gender to the author. If there was a discrepancy across components, we identified the most popular of the components and assigned that gender to the author. We cross checked the quality of our gender assignment algorithm manually.

11.5.4 Computing citation counts of authors

We directly pulled estimated citation counts for each paper from the *Microsoft Academic Graph* (MAG) data set and from the *Crossref* dataset when the MAG information was unavailable. We then assigned to each author the sum of citations of the author's articles.

11.5.5 Assignment of institutional affiliations of authors

For a subset 47 US institutions, we matched the theorists in our sample with their home department by using a dataset a manually collectedby us. A department is included if it is in the top 50 list of the *RePEc* US department rankings in 2013, 2014 and 2015.[42] The department level dataset covered all faculty members as well as their titles from 1995 to 2019 from two sources (department websites and course catalogues). We matched our sample of theorists to the faculty members in these departments using their names. This sums up to a total of 11,087 theorists with affiliation info.

11.6 Description of the methodology to estimate the community detection model based on Feng et al. (2023)

Taking logs from (6), we can express the log likelihood compactly as

$$\log \mathcal{L} = \sum_{t\in\{\ell,c\}} n_t(\boldsymbol{\tau}) \log(\pi_t) + \sum_{t,t'\in\{\ell,c\}} M_{tt'}(\boldsymbol{\tau}) \log(\omega_{tt'}) - \sum_{t,t'\in\{\ell,c\}} \omega_{tt'} B_{tt'}(\boldsymbol{\tau},\boldsymbol{\gamma}) + \sum_{i=1}^{n}\sum_{j=i}^{n} q_{ij} a_{ij} \mathbf{x}_{ij}'\boldsymbol{\gamma} \tag{9}$$

[41] See https://www.rdocumentation.org/packages/genderizeR/versions/2.0.0.
[42] See https://ideas.repec.org/top/old/1505/top.usecondept.html.

where $q_{ij} = 1$ if $j \in Q(i)$,

$$n_t(\boldsymbol{\tau}) = \sum_{i=1}^{n} 1\{\tau_i = t\}$$

is the total number of type t authors under assignment $\boldsymbol{\tau}$,

$$M_{tt'}(\boldsymbol{\tau}) = \sum_{i=1}^{n}\sum_{j=i}^{n} q_{ij}a_{ij}1\{\tau_i = t, \tau_j = t'\}$$

is the number of co-authorships between a type t and a type t' authors under assignment $\boldsymbol{\tau}$, and

$$B_{tt'}(\boldsymbol{\tau}, \boldsymbol{\gamma}) = \sum_{i=1}^{n}\sum_{j=i}^{n} q_{ij}e^{\mathbf{x}'_{ij}\boldsymbol{\gamma}}1\{\tau_i = t, \tau_j = t'\}$$

is an aggregate of the covariate influence in co-authorship formation among type t and a type t' authors under assignment $\boldsymbol{\tau}$.

We can first take the FOC with respect to π_t and Ω. With respect to π_ℓ:

$$\begin{aligned} & n_\ell(\boldsymbol{\tau})\frac{1}{\pi_\ell} + (n - n_\ell(\boldsymbol{\tau}))\frac{1}{1-\pi_\ell}(-1) = 0 \\ & \Rightarrow \\ & \pi_\ell^{MLE} = \frac{n_\ell(\boldsymbol{\tau})}{n} \end{aligned} \tag{10}$$

With respect to $\omega_{tt'}$,

$$\begin{aligned} & \frac{M_{tt'}(\boldsymbol{\tau})}{\omega_{tt'}} - B_{tt'}(\boldsymbol{\tau}, \boldsymbol{\gamma}) = 0 \\ & \Rightarrow \\ & \omega_{tt'}^{MLE} = \frac{M_{tt'}(\boldsymbol{\tau})}{B_{tt'}(\boldsymbol{\tau}, \boldsymbol{\gamma})} \end{aligned} \tag{11}$$

Plugging back (10) and (11) into (9), we obtain the profile likelihood:

$$\begin{aligned} \log \mathcal{L}^* = & \sum_{t\in\{\ell,c\}} n_t(\boldsymbol{\tau}) \log\left(\frac{n_t(\boldsymbol{\tau})}{n}\right) + \sum_{t,t'\in\{\ell,c\}} M_{tt'}(\boldsymbol{\tau}) \log\left(\frac{M_{tt'}(\boldsymbol{\tau})}{B_{tt'}(\boldsymbol{\tau}, \boldsymbol{\gamma})}\right) \\ & - \sum_{t,t'\in\{\ell,c\}} \frac{M_{tt'}(\boldsymbol{\tau})}{B_{tt'}(\boldsymbol{\tau}, \boldsymbol{\gamma})} B_{tt'}(\boldsymbol{\tau}, \boldsymbol{\gamma}) + \sum_{i=1}^{n}\sum_{j=i}^{n} q_{ij}a_{ij}\mathbf{x}'_{ij}\boldsymbol{\gamma} \end{aligned}$$

Notice that the third sum is a constant equal to the total number of co-authorships, so it does not depend on $\boldsymbol{\tau}$ or $\boldsymbol{\gamma}$.

Thus, maximizing $\log \mathcal{L}^*$ is equivalent to maximizing

$$\log \tilde{\mathcal{L}}^* = \sum_{t \in \{\ell, c\}} n_t(\boldsymbol{\tau}) \log\left(\frac{n_t(\boldsymbol{\tau})}{n}\right) + \sum_{t,t' \in \{\ell,c\}} M_{tt'}(\boldsymbol{\tau}) \log\left(\frac{M_{tt'}(\boldsymbol{\tau})}{B_{tt'}(\boldsymbol{\tau}, \boldsymbol{\gamma})}\right) + \sum_{i=1}^{n}\sum_{j=i}^{n} q_{ij} a_{ij} \mathbf{x}_{ij}' \boldsymbol{\gamma} \tag{12}$$

For a given ideological assignment $\tilde{\boldsymbol{\tau}}$, the terms of the form $n_t \log(n_t/n)$ and $M_{tt'} \log(M_{tt'})$ do not depend on $\boldsymbol{\gamma}$, so

$$\hat{\boldsymbol{\gamma}}(\tilde{\boldsymbol{\tau}}) = \text{argmax}_{\,\gamma} \left\{ \sum_{i=1}^{n}\sum_{j=i}^{n} q_{ij} a_{ij} \mathbf{x}_{ij}' \boldsymbol{\gamma} - \sum_{tt' \in \{\ell,c\}} M_{tt'}(\tilde{\boldsymbol{\tau}}) \log\left(B_{tt'}(\tilde{\boldsymbol{\tau}}, \boldsymbol{\gamma})\right) \right\}$$

This objective function is strictly concave in $\boldsymbol{\gamma}$, so it has a unique solution that can be easily found with a BFGS algorithm.

We can now plug in $\hat{\boldsymbol{\gamma}}(\tilde{\boldsymbol{\tau}})$ in (7):

$$\log \tilde{\mathcal{L}}^*(\tilde{\boldsymbol{\tau}}) = \sum_{t \in \{\ell, c\}} n_t(\tilde{\boldsymbol{\tau}}) \log\left(\frac{n_t(\tilde{\boldsymbol{\tau}})}{n}\right) + \sum_{t,t' \in \{\ell,c\}} M_{tt'}(\tilde{\boldsymbol{\tau}}) \log\left(\frac{M_{tt'}(\tilde{\boldsymbol{\tau}})}{B_{tt'}(\tilde{\boldsymbol{\tau}}, \hat{\boldsymbol{\gamma}}(\tilde{\boldsymbol{\tau}}))}\right) + \sum_{i=1}^{n}\sum_{j=i}^{n} q_{ij} a_{ij} \mathbf{x}_{ij}' \hat{\boldsymbol{\gamma}}(\tilde{\boldsymbol{\tau}}) \tag{13}$$

The space of possible vectors $\boldsymbol{\tau}$ is very large; there are 2^n possible vectors. Feng et al. (2023) propose an algorithm that works very well:

1. Pick an arbitrary $\tilde{\boldsymbol{\tau}}$, and find $\hat{\boldsymbol{\gamma}}(\tilde{\boldsymbol{\tau}})$.
2. Maximize (13) evaluated at $\hat{\boldsymbol{\gamma}}$ using an EM algorithm. For details on the EM algorithm, see Feng et al. (2023).
3. This yields an allocation $\tilde{\boldsymbol{\tau}}(\tilde{\boldsymbol{\gamma}})$.
4. Iterate if desired, although in practice the first iteration will already deliver a very accurate allocation.

11.7 Description of the methodology to estimate the multinomial choice model through simulated maximum likelihood

We maximize (8) using the method of maximum simulated likelihood. This entails numerically simulating the double integral that averages over the distribution of peer effects conditional on psychological types, and then averaging over those types. We simulate this

integral with a discrete sum. The estimator takes the form

$$ln\hat{L}(\boldsymbol{\gamma}) = \sum_{a=1}^{N} \sum_{\rho \in \{m,f,x,p\}} 1\{p_{a(ij)t} = \rho\} \times$$
$$ln\left[\frac{1}{B_1}\frac{1}{B_2}\sum_{b_i=1}^{B_1}\sum_{b_j=1}^{B_2}\sum_{\psi_i,\psi_j \in \{\underline{\psi},\overline{\psi}\}} G_\rho\left(V^{\rho}_{a(ij)t}(b_{i|\psi_i}, b_{j|\psi_j})\right)\mathbb{P}(\psi_i|\mathbf{w}_i)\mathbb{P}(\psi_j|\mathbf{w}_j)\right],$$

where

$$G_\rho(v^\rho) = \frac{\exp(v^\rho)}{1 + \sum_{s\in\{m,f,x\}}\exp(v^s)},$$

and the b_i, b_j are draws for the β coefficients for each author from Beta distributions, and B_1, B_2 are the number of draws for approximating the integrals. For single-authored papers the integral is one dimensional.