Journal of Economic Behavior and Organization 238 (2025) 107165

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Research paper

The gender gap in preferences: Evidence from 45,397 Facebook interests[☆]

Ángel Cuevas [a,b], Rubén Cuevas [a,b,c], Klaus Desmet [d,e,f,*], Ignacio Ortuño Ortín [g]

[a] Department of Telematic Engineering, Universidad Carlos III de Madrid, Spain
[b] Hiili S.L., Spain
[c] UC3M-Santander Big Data Institute, Spain
[d] Department of Economics and Cox School of Business, Southern Methodist University, United States of America
[e] NBER, United States of America
[f] CEPR, United Kingdom
[g] Department of Economics, Universidad Carlos III de Madrid, Spain

ARTICLE INFO

JEL classification:
D91
J16
Z1

Keywords:
Gender gap
Gendered preferences
Non-gendered preferences
Facebook
Gender equality paradox

ABSTRACT

This paper uses information on the frequency of 45,397 Facebook interests to study how the difference in revealed preferences between men and women changes with a country's degree of gender equality. For preference dimensions that are systematically biased toward the same gender across the globe, differences between men and women are larger in more gender-equal countries. In contrast, for preference dimensions with a gender bias that varies across countries, the opposite holds. The so-called gender-equality paradox is therefore limited to preferences that display the same systematic gender bias across societies.

1. Introduction

It is well established that for certain preferences – such as those related to educational choices or competitiveness – differences between men and women tend to be larger in countries with higher levels of gender equality. This paper asks to what extent this gender-equality paradox generalizes to the near-universe of preferences, using evidence from the prevalence of 45,397 Facebook interests by gender. We identify two groups of interests that exhibit contrasting patterns. For interests that are systematically biased toward the same gender across countries, the gender-equality paradox holds: differences between men and women are larger in more gender-equal societies. In contrast, for interests that do not display a systematic gender bias, the opposite pattern emerges: men and women are more alike in countries with greater gender equality.

Our data on the frequency of interests by gender and country come from Facebook. The social media company observes each of its almost three billion users' online activity, not just on its own platform, but also on all websites and apps where it has a presence.

☆ We thank Iñaki Ayensa for his help with the Facebook data, and we are grateful to Lola Collado, David Cuberes, Cynthia Desmet Villar, Ömer Özak, Javier Ruiz-Castillo, Romain Wacziarg and Beatriz Yankelevich for helpful comments. The authors furthermore acknowledge the financial support of Fundación BBVA (Ayudas a equipos de investigación for project "Aplicaciones económicas de distancias culturales"); MICIN/AEI grants PID2019-109157RB-IOO and CEX2021-001181-M; MCIN and the NextGeneration EU/PTR Fund grant TED2021-130118B-I00 for the project ENTRUDIT; Comunidad de Madrid, grant EPUC3M11 (V PRICIT); Ministry of Economic Affairs and Digital Transformation and the EU NextGeneration Fund (PRITIA-CLOUD project); and EU FEDER Fund.

* Corresponding author at: Hiili S.L., Spain.
E-mail addresses: acrumin@it.uc3m.es (Á. Cuevas), rcuevas@it.uc3m.es (R. Cuevas), kdesmet@smu.edu (K. Desmet), iortuno@uc3m.es (I.O. Ortín).

https://doi.org/10.1016/j.jebo.2025.107165
Received 22 January 2024; Received in revised form 26 June 2025; Accepted 18 July 2025

Available online 2 September 2025
0167-2681/

In addition, it tracks many of its users' offline activities through GPS tracking. Facebook uses this information to assign interests to all of its users. In doing so, it has arguably created the world's largest database on human interests. At its essence, Facebook infers its users' interests and preferences by observing their activity. The easiest way to understand this process is by relying on a revealed preference argument.

By querying Facebook's publicly available Marketing API, we collect for most countries of the world the number of male and female users interested in 45,397 different topics. Because the data are at the level of populations (e.g., Canadian men or Ghanaian women), they do not involve any individual privacy issues. Compared to other potential data sources on preferences, Facebook data have two key advantages. First, it covers the near-universe of interests, ranging from religious beliefs and sports, to political preferences and cuisine. Second, in contrast to surveys, Facebook interests constitute a bottom-up revealed measure of preferences, covering whatever users find interesting, rather than what social scientists deem important.

We start by examining whether the gender-equality paradox exists for the full set of 45,397 interests. For each country, we compute the cosine distance between the interest frequency vectors of men and women, providing a country-level measure of the overall difference in interests between genders. We then analyze how this measure is related to the degree of "vertical" gender equality, defined as the absence of a male-primacy ideology that perpetuates male dominance in different realms of society.[1] To measure vertical gender equality, we rely on a standard index that captures equality in access of men and women to resources and opportunities. We find a weak, positive association between gender equality and the interest gap between men and women, but this relationship is not robust to the inclusion of different controls. Hence, when considering all interests jointly, our findings on the gender-equality paradox are inconclusive.

Next, we differentiate between gendered and non-gendered interests. We say that an interest is gendered if it displays a systematic bias toward the same gender across the globe. More specifically, if in more than 90% of countries an interest is more prevalent among the same gender, then we refer to it as gendered. For example, "cosmetics" and "motherhood" are universally more common among women, whereas "motorcycles" and "Lionel Messi" are universally more common among men. Conversely, we say that an interest is non-gendered if its gender bias varies across countries. More specifically, if an interest is more common among men in at least 30% of countries and more common among women in at least another 30% of countries, then we refer to it as non-gendered. For example, "world heritage site" and "physical fitness" do not display a systematic gender bias across the globe.

When exploring the relationship between a country's gender equality and differences in interests between men and women, we uncover a sharp distinction between gendered and non-gendered interests. Greater gender equality is associated with larger differences between men and women for gendered interests, whereas the opposite holds for non-gendered interests. Thus, the gender-equality paradox generalizes to the set of gendered interests, but there is no such paradox for non-gendered interests.

While using thousands of interests has the advantage of being comprehensive and inclusive, it may also generate noise that leads to biases. For example, if one group is interested in "fatherhood" and another in "motherhood", we might overestimate differences if both reflect a common interest in "parenting". Because different interests may sometimes represent the same underlying preferences, we assess the robustness of our findings by applying singular value decomposition to the data matrix, allowing us to identify the main latent preference dimensions. Another potential bias relates to users being reluctant to reveal certain interests. However, because users need not explicitly declare their interests for Facebook's algorithm to identify them, this is unlikely to be an important concern. Moreover, even if such reluctance were to introduce noise, it would be attenuated by singular value decomposition, which leverages the entire correlation structure between all interests to identify the main latent preference dimensions.

Once we have identified the main preference dimensions through singular value decomposition, we can again distinguish between gendered and non-gendered dimensions of preferences. For a preference dimension to be gendered, we require the relative positions of men and women along that dimension to be similar across countries. With this alternative method, we confirm the paper's central result: more gender-equal societies tend to be associated with greater differences between men and women in gendered preferences but smaller differences in non-gendered preferences.

This paper relates to a vast literature in psychology, sociology and economics that examines whether differences in values, attitudes, interests and personality get accentuated or attenuated in societies with higher levels of gender equality. Balducci (2023) and Herlitz et al. (2025) provide excellent meta-analyses of this question. Most empirical studies in this area have focused on gender differences in personality characteristics (Costa et al., 2001; Kaiser, 2019; Mac Giolla and Kajonius, 2019), cognitive abilities (Lippa et al., 2010), educational preferences (Stoet and Geary, 2018), basic human values (Fors Connolly et al., 2020), and specific cultural, behavioral and moral values (Falk and Hermle, 2018; Atari et al., 2020). Many of these papers find evidence of the gender-equality paradox. For example, more gender-equal countries are found to exhibit greater sex differences in care and fairness (Atari et al., 2020), altruism, trust and risk-taking (Falk and Hermle, 2018), willingness to compete (Cárdenas et al., 2012), and the big five personality traits (Mac Giolla and Kajonius, 2019). However, some other studies find the opposite or argue that the relationship is not robust. For example, Guiso et al. (2008) show that the math gender gap narrows in societies with greater gender equality,

Our paper differs from this previous work in three respects. First, most studies have focused on particular traits, values, preferences, or abilities. For example, Stoet and Geary (2018) examine educational preferences for STEM fields, while Falk and Hermle (2018) explore five specific aspects of preferences (willingness to take risks, patience, altruism, positive and negative reciprocity, and trust). This raises the question of whether the gender-equality paradox generalizes to other preferences and interests. By analyzing comprehensive data on 45,397 interests, we are able to draw systematic and general conclusions about gender

[1] This is distinct from the notion of "horizontal" gender equality, i.e., the belief that men and women are different, without one dominating the other. We return to this discussion later in the Introduction.

differences in preferences. Second, to study the relationship between gender equality and preferences, a broad cross-section of countries is essential. Our baseline study includes 98 countries, and we also report results for a sample of 110 countries. By contrast, most of the cross-country studies cited above cover 20 to 50 countries, with Falk and Hermle (2018) being an exception, using a sample of 76 countries. Third, papers that examine the effect of gender equality on differences in preferences often overlook the possibility of reverse causality. As part of our robustness analysis, we address this endogeneity concern by taking an instrumental variable approach. Our results are suggestive of a causal interpretation of the paper's main finding.

Also related to our work is the literature that seeks to identify some of the key differences in preferences between men and women. Many experimental studies have documented systematic gender differences in risk attitudes, dislike of competition, and social preferences (see (Croson and Gneezy, 2009; Bertrand, 2011; Niederle and Vesterlund, 2011), for excellent surveys).[2] An important, related, question is the extent to which these differences are due to nature or nurture. Evidence for the role of nature comes from studies showing that hormones may affect certain preferences, such as attitudes toward competition and risk-taking, as well as career choices and activities (Archer, 2006; Sapienza et al., 2009; Buser, 2012; Wozniak et al., 2014; Berenbaum and Beltz, 2021). However, a recent large-scale study of 3450 individuals finds no statistically significant association between prenatal testosterone exposure and economic preferences (Neyse et al., 2021). This is consistent with older work by Gneezy et al. (2008) who show that gender differences in risk-taking are society-dependent, ruling out a purely nature-based explanation. Other papers that have shown a role for culture and environment in shaping gender differences include Aldén and Neuman (2022) and Holmlund et al. (2023) who examine the gender gap in educational choices among second-generation immigrants, as well as Molina and Usui (2023) who find that gender gaps in educational attainment and marriage are smaller in areas with higher female employment.

Our paper contributes to a growing body of research that leverages Facebook's Marketing API to study a variety of socioeconomic issues, such as immigrant assimilation (Dubois et al., 2018), political campaigns (Liberini et al., 2020) and crime rates (Fatehkia et al., 2019). Facebook ads data have also been used to investigate specific aspects of the gender gap. For instance, Vieira and Vasconcelos (2021) analyze the gender balance in STEM in Brazil; Carrer and De Masi (2024) use Facebook data to create a gender norms measure at the municipal level in Italy; Mejova et al. (2018) study the digital gender gap in India; and García et al. (2018) show how gender inequality in Facebook usage correlates with various aspects of gender inequality. With a few exceptions, however, most papers rely on relatively small sets of Facebook interests. One exception is Dubois et al. (2018) who includes around 3000 Facebook interests, but it is limited to one country. Another is some of our own recent work, where we use 60,000 interests to measure cultural differences between countries (Obradovich et al., 2022). By correlating these differences with traditional survey-based measures of cross-cultural distances, that paper provides proof of concept for using social media data to measure cultural and preference differences across populations. The methodology established there justifies our use of Facebook interests to measure preference differences between men and women.

The rest of the paper is organized as follows. Section 2 describes the data, with a special emphasis on the Facebook data on interests; Section 3 analyzes the relation between gender equality and gender differences in interests; Section 4 explores how this relation depends on whether interests are gendered or not; Section 5 uses single value decomposition to identify the main gendered and non-gendered preference dimensions; Section 6 discusses the findings; and Section 7 concludes.

2. Data

This paper asks whether greater gender equality amplifies or attenuates differences in interests between men and women. Because we want to treat this question comprehensively, our biggest challenge is to get data on the prevalence of many different interests by gender for a large cross-section of countries. Below we describe how we achieve this by obtaining information on the frequency of 45,397 Facebook interests by gender and country. We also discuss, more briefly, the data on gender equality and other control variables. Online Appendix B describes the data sources in more detail.

2.1. Dependent variable: Gender differences in interests

Data on interests by gender and country. Of its almost 3 billion Facebook users worldwide, Facebook observe their likes, shares, clicks and downloads, not only on its own platform but also on all other websites and apps where the company is present.[3] Additionally, through access to GPS data, Facebook also observes some of its users' offline activity. By unobtrusively observing its users, the social media company has created a massive database on people's revealed preferences and interests. We refer to these interests as "revealed" in the sense that they are identified by observing users' actual behavior.[4] This approach enables us to measure interests from the bottom up, without selectively excluding or including specific interests. It stands in contrast to the traditional top-down approach used in surveys or experiments, where researchers decide which few select features to investigate. In our case, we include everything users reveal an interested in, offering a more comprehensive and less biased view of people's interests.

Facebook does not directly give us a list of available interests. Instead, we need to prompt Facebook's Marketing API to provide us with a comprehensive and broad set of interests. To do so, we take the 1000 most common words in English, as well as all

[2] See also Samek (2019) and Abraham (2023) for more recent evidence.

[3] According to Englehardt and Narayanan (2016), Facebook has a presence in over 30% of the 1 million most popular websites.

[4] Appendix A provides a simple revealed preference framework that helps to understand the connection between the interests that Facebook assigns to its users and their preferences. Essentially, the idea is that Facebook receives signals about its users' preferences by observing the time they spent on different topics and subjects.

possible combinations of one, two and three letters (‘a’, ‘b’, ..., ‘aa’, ‘ab’, ..., ‘aaa’, ‘aab’, ...). By not using a predefined list of interests, we capture a comprehensive set of possibilities. Indeed, for each one of these words and letter combinations, we query the Facebook Marketing API for up to 1,000 interests that match or contain these words or letters. This generates a list of 308,568 unique interests. All interests come with a numerical code that is common to all languages and users. Although this may not cover the entire universe of interests, it is an extensive list.

Of the 308,568 interests we identify, we retain those with a worldwide Facebook audience of more than one million but less than one billion. This yields the 45,397 interests that we use in this paper. We apply a minimum threshold of one million for two reasons. First, to keep the task of getting information on users’ interests by gender and country manageable, it is necessary to reduce the total number of interests. Second, Caravaca et al. (2024) show that Facebook is more accurate in identifying the more popular interests, so it makes sense to focus on interests that exceed this minimum threshold. We introduce a maximum threshold of one billion, because the 39 interests that surpass this limit have an outsized effect on our measure of the distance between men and women. When excluding these interests, the distances between men and women based on our 45,397 interests become highly robust. For instance, the distances are almost identical to the ones obtained from a random sample of 25,000 interests. Additionally, these 39 interests tend to be very generic (e.g., “Facebook”, “social networks”) and are assigned to almost everyone, so we do not want them to distort our distance measure.

For each one of the 45,397 interests, we query Facebook’s Marketing API for the corresponding number of monthly active users (MAU) by gender and country. This gives us, for example, the number of female users in France interested in “Youssou N’Dour” or the number of male users in Singapore interested in “chili crab”. To automate the querying process, we developed a Facebook audience capture and analysis tool. Even when automated, this is a lengthy and time-consuming effort that spanned the entire first semester of 2019.

The Facebook data we use do not raise any privacy concerns. The Marketing API gives us information at the level of population groups, and never at the level of individuals. To further ensure anonymity, the minimum number of monthly active users (MAU) reported by the API for any demographic is 1,000. While in principle this can distort our distance measure between men and women, this is not an issue as long as the number of interests is large enough and as long as groups are not too small.[5]

Measuring gender differences in interests. Let f^w_{ci} be the number of female Facebook users in country c who hold interest i, and let f^m_{ci} be the corresponding number of male users. We can then write the vector with the interest frequencies of women in country c as $f^w_c = \{f^w_{c1}, f^w_{c2}, \dots, f^w_{cI}\}$. The corresponding vector for men is f^m_c. For example, an element of f^w_{ci} could be the total number of female Facebook users of country c who have an interest in “beer”.

The cosine distance between the interest vector of men and women in country c measures the gender difference in interests in that country:[6,7]

$$CosDist_c = 1 - \frac{\sum_{i=1}^{I} f^m_{ci}\, f^w_{ci}}{\sqrt{\sum_{i=1}^{I}(f^m_{ci})^2}\sqrt{\sum_{i=1}^{I}(f^w_{ci})^2}}. \tag{1}$$

We can compute the difference between men and women using all interests or a subset of interests. Later in the paper, we distinguish between gendered and non-gendered interests, and compute the cosine distance between men and women for each one of the subsets.

In Obradovich et al. (2022) we show that this distance in Facebook interests is a good measure of cultural distances between populations. The average within-country distance between men and women is 0.08. To put this number in context, the average distance between populations of different countries is 0.25. Fig. 1 depicts the gender differences in Facebook interests in 149 countries with population above one million, Facebook penetration rate above 2.5% and number of Facebook users greater than 100,000. Online Appendix Table C.1 provides the full data.

2.2. Main independent variable: Gender equality

Most work on the gender-equality paradox focuses on the impact of “vertical” gender equality—the degree of equality in rights and opportunities between men and women. We follow this approach and take the 2018 World Economic Forum’s Gender Gap Index (WEF) as our main measure of gender equality. This index is one of the best-established indices of vertical gender equality. It measures gender-based gaps in access to resources and opportunities in countries, rather than the actual level of those countries’ available resources and opportunities. Thus, this index captures the level of gender equality separately from the level of economic development. The index is increasing in the degree of equality and has a scale from zero to one. It is made up of four subindices, related to economic opportunity, educational attainment, health outcomes and political empowerment. Examples of variables that contribute to the WEF gender equality index include female labor force participation relative to male, female earned income over male, sex ratio at birth, gender difference in healthy life expectancy, and females with seats in parliament. Fig. 2 shows a world map of the WEF gender equality index for the same sample of countries as Fig. 1.

[5] More specifically, if we choose a subset of most popular interests, our distance measure is virtually unchanged for any threshold above 20,000 interests. If, instead, we use a random subset of interests, the corresponding threshold is around 25,000 interests. When the number of interests is small, Rama et al. (2020) offer another solution to alleviate this problem.

[6] The cosine distance is the most commonly used distance metric if the data are high-dimensional. It has the additional advantage that it does not depend on how the data are normalized. Note that the cosine distance is proportional to the Euclidean distance if we normalize all the vectors to have modulo 1.

[7] Note that the cosine distance does not change if we use interest shares.

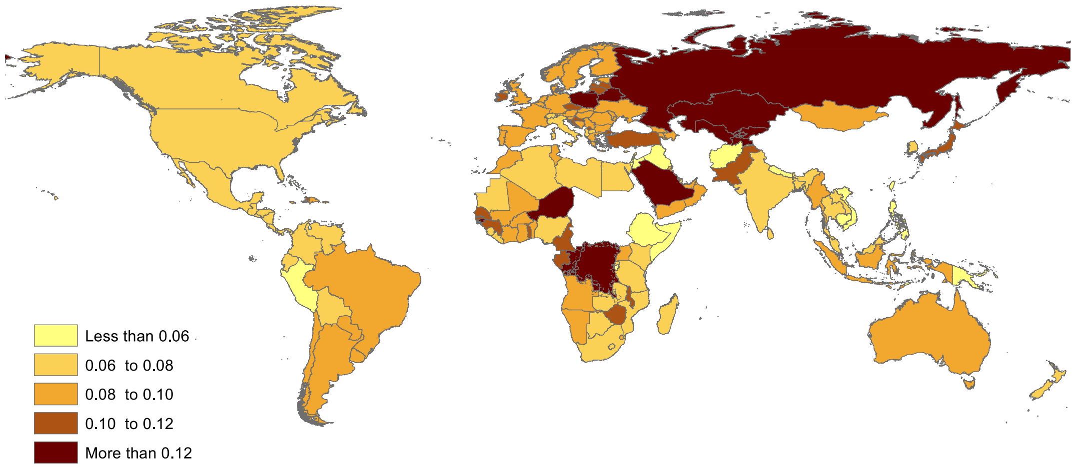

Fig. 1. Gender Differences in Facebook interests.
Figure shows the cosine distance between the interest frequency vector of men and women, based on 45,397 Facebook interests in countries with population above one million, Facebook penetration rate above 2.5% and number of Facebook users greater than 100,000.

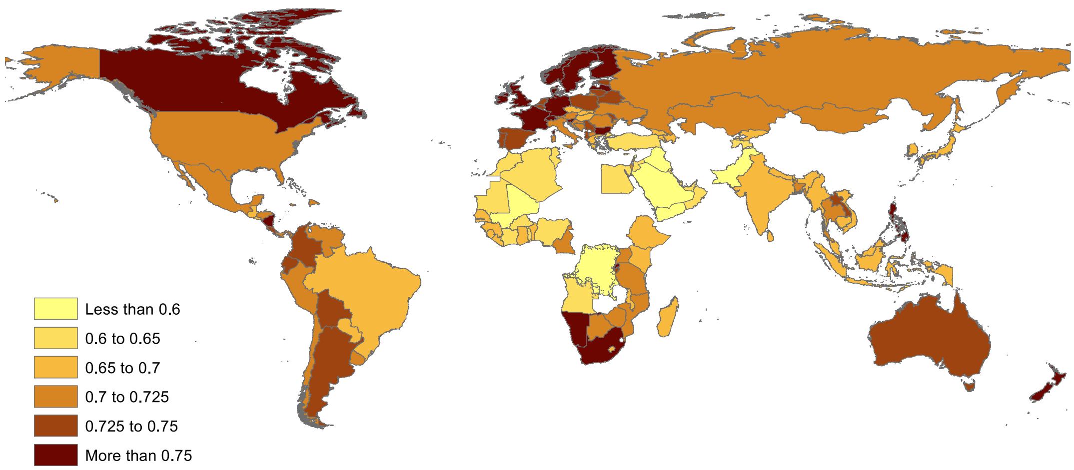

Fig. 2. Gender equality.
Figure depicts the 2018 Gender Gap Index of the World Economic Forum (WEF). The index is increasing in gender equality, and is based on a set of metrics related to economic opportunity, educational attainment, health outcomes and political empowerment. Sample of countries is the same as that in Fig. 1, with the exception of a few countries for which the Gender Gap Index is not available.

For robustness purposes, we also consider alternative indices, such as the UNDP's Gender Inequality Index and the OECD's Social Institutions and Gender Index. UNDP's Gender Inequality Index is similar to the WEF index: it measures inequality in reproductive health, educational attainment, political empowerment and economic status. As for the OECD's Social Institutions and Gender Index, it aims to capture discrimination against women in formal and informal social institutions. More specifically, it measures discrimination in the family, restricted physical integrity, restricted access to productive and financial resources, and restricted civil liberties. For instance, it includes measures of discrimination in divorce and inheritance laws, violence against women, genital mutilation, workplace rights, and access to justice and financial services. For reasons of comparison with the WEF index, we recode the UNDP and the OECD indices so that both are increasing in the degree of gender equality.

2.3. Other control variables

Other variables are likely to affect a country's gender differences in interests. In our baseline specification we include two additional control variables: the level of economic development and the overall diversity in interests. In other specifications, we add further controls.

Economic development. An increase in income per capita reduces material constraints, allowing men and women to more freely express gender-specific desires, interests and ambitions (Falk and Hermle, 2018). Because economic development and gender equality are positively correlated (Fernández, 2014; Cuberes and Teignier, 2014), it is important to control for the separate effect of income per capita on preference differences. In our baseline sample, the correlation between income per capita and gender equality is 0.21. While positive, the correlation is far from perfect: there are poor countries with a high degree of gender equality, like Uganda, and rich countries with a low degree of gender equality, like Saudi Arabia. While we use income per capita as our main measure of development, we also explore whether our results are robust to using the human development index (HDI).

Overall diversity. Bigger differences between men and women could partly reflect greater overall heterogeneity in society. Or on the contrary, more pluralistic countries in terms of interests might display smaller gender differences. To control for a country's overall diversity, we use the entropy index, given by $Ent_c = -\sum_{i=1}^{I} s_{ci} \log(s_{ci})$, where $s_{ci} = f_{ci} / \sum_i f_{ci}$ and f_{ci} is the number of individuals in country i who hold interest i. In the baseline sample of countries, the correlation between $CosDist_c$ and Ent_c is -0.25, implying that on average countries with greater gender differences in interests exhibit less overall heterogeneity in interests.

Other controls. In our robustness checks, we also consider more comprehensive specifications where we control for additional variables. First, we include regional dummies. This allows us to evaluate whether our results are mostly driven by differences between the world's large regions, or whether they also hold within regions. Second, we assess the importance of a country's religious composition and its exposure to Soviet influence. Religious beliefs affect gender norms and roles, and may hence be a confounding factor. Soviet influence may also have shaped gender norms, since official Soviet doctrine viewed gender differences to be irrelevant. Third, we control for the degree of Facebook penetration. Data from countries with low penetration rates may be less representative and reliable. Fourth, we check for the possible role of geographic and climatic variables in shaping gender preferences. For example, geo-climatic conditions may have affected the division of labor between men and women, leading to specific gender norms that have persisted.

2.4. Further discussion of Facebook data

In this subsection we provide some further discussion on the use of Facebook data. We address a number of potential concerns and argue that our methodology provides a reliable way to measure the interests and preferences of population groups.

Facebook's incentive to correctly identify interests. Why would Facebook have an incentive to correctly identify its users' interests and preferences? And why would these include interests that are not directly related to marketable goods and services? At its core, Facebook's business model relies on monetizing the time users spend on the social network, as that time directly determines how many ads it can show them. To keep its users engaged, Facebook must show them information that is closely related to their true preferences and interests.

These include not just goods and services that can be bought on the marketplace, but also any non-market activity or interest that gives users utility. For example, if Facebook identifies a user who likes mountain landscapes, it can keep that user engaged by displaying pictures of mountains. This gives Facebook more time to show her targeted ads that may very well be completely unrelated to mountains. Because of this, Facebook has an incentive to learn its users' preferences, including those that do not refer to marketable goods and services. Moreover, even if there were a subset of interests that were of no use at all to Facebook, its algorithm would likely still identify them. Indeed, by tracking an individual's overall activity, it gets a comprehensive view of her interests and preferences. In that sense, Facebook acts as an accidental ethnographer, observing all of its users' activities, without judging what is more relevant or less relevant (Obradovich et al., 2022).

That might still leave the question why Facebook's Marketing API would make interests that are not marketable available to potential advertisers. Here the answer is more straightforward: although we would probably not view an interest such as "fatherhood" or "God" as being a marketable good or service, advertisers might still want to target their publicity to users for whom "fatherhood" or "God" are important.

How well does Facebook identify interests? Having an incentive to correctly identify interests is not the same as being successful at doing so. What is the evidence on Facebook's effectiveness?

First, a survey of 963 U.S. Facebook users conducted by the Pew Research Center (2019) found that only 27% felt the interest categories assigned to them by Facebook were inaccurate. Among those assigned a political category, 73% believed Facebook's categorization of their political views was very or somewhat accurate.

Second, other studies have shown that the digital footprint of Facebook users can accurately predict personal traits. Based on around 70,000 participants, Youyou et al. (2015) show that Facebook Likes more accurately predict personality traits than friends and family—and as accurately as one's spouse. In a related study involving 58,000 participants, Kosinski et al. (2013) find that Facebook Likes can accurately predict sexual orientation, ethnicity, religious views, political views, intelligence and happiness.

Notably, these two studies rely only on Facebook Likes, whereas Facebook's actual algorithm considers a user's entire digital footprint, so we would expect its accuracy in predicting interests and preferences to be higher.

Third, there is even stronger evidence that Facebook is highly effective at identifying interests and preferences at the level of population groups, rather than at the level of individuals. Obradovich et al. (2022) provides proof of concept, demonstrating that Facebook interests can be used to measure preference differences between countries, regions, and population groups. The paper establishes that distances between countries based on 60,000 Facebook interests display small positive correlations with linguistic, geographical, religious and genetic distances, but a substantially higher correlation of 0.54 with a more direct measure of cultural and preference differences based on the World Values Survey. It also applies a community detection algorithm to Germany and India, revealing two communities in Germany aligned with the historical east–west divide and three communities in India that roughly correspond to the country's linguistic regions.

Obradovich et al. (2022) furthermore demonstrate that hierarchical clustering of U.S. states based on Facebook distances produce logical groupings: Midwestern states cluster together, as do Southern states. Mountainous and more rural states also appear together, with Alaska being closest to states like North Dakota, Idaho and New Hampshire, despite the significant geographical distances between them. Using the same hierarchical clustering technique, Online Appendix Figure C.1 displays a similar dendrogram for countries using our data. Groupings are largely as expected: culturally or historically connected countries are clustered together (e.g., Germany and Austria, Brazil and Portugal, or the U.S., Canada, Ireland and the U.K.). Overall, we conclude that Facebook data on interests provides a good way to measure preference differences between population groups.

Users' revelation of interests. Might some users be reluctant to reveal certain interests, hence biasing our measure? This is unlikely to be a major concern. First, recall that Facebook's algorithm does not rely on users explicitly declaring their interests. Instead, it identifies users' interests based on their overall online activity as well as on part of their offline activity through GPS tracking. This makes it hard for a user to fool Facebook's algorithm. As an example of this difficulty, Cabañas et al. (2020) prove that Facebook labeled with the interest "homosexuality" hundreds of thousands of users in countries where homosexuality is severely punished (Saudi Arabia, Somalia, Qatar, UAE, etc.). In these countries users are unlikely to proactively declare an interest in homosexuality, and yet the interest was inferred by Facebook's algorithm.

Second, our study is based on over 45,000 different interests. Even if there are a few topics that a user might try to avoid, this is unlikely to bias our measures in a significant way. That said, it may make our measures more noisy. To address this concern, we use singular value decomposition, a technique similar to principal component analysis that identifies the main latent preference dimensions. By taking into account the entire correlation structure between all interests, this methodology allows to correctly identify preferences even if a few select interests are avoided.

Third, Facebook users' revelation of interests is less liable to biases that plague people's revelation of preferences through surveys. It is well known that surveys pose substantial risk of social desirability bias and Hawthorne effects (i.e., the alteration of behavior due to being observed). In contrast, because Facebook observes its users unobtrusively, their behavior is less likely to be influenced by such biases.

Representativity of sample of Facebook users. To ensure that Facebook users are sufficiently representative of the population groups that we are interested in, our baseline sample consists of 106 countries with a population above one million and a Facebook penetration rate of at least 25%. In our robustness checks, we consider an alternative sample that lowers the Facebook penetration thresholds to 5%.[8] Reassuringly, the main results are unchanged. The gender composition of our sample is quite balanced. The average share of women across countries is 0.47, with that share being 0.50 in countries where gender equality is above the median and 0.43 in countries where gender equality is below the median.

Another issue affecting cross-country comparability is that Facebook users may be more biased toward younger populations in some countries than in others. The share of young, defined as those aged 40 and below, is 0.67 in our sample. That proportion is 0.62 in countries with gender equality below the median, and 0.71 in countries with gender equality above the median. To test the robustness of our results to age composition, we control for the ratio of young to old Facebook users. In addition, we obtain the interest frequencies by age and gender for a random subsample of 5000 interests, and re-run our main regressions separately for the old and the young.

Synonyms and language. While using tens of thousands of interests has many advantages, it may generate noise that leads to systematic biases. For example, if one group is interested in "fatherhood" and another in "motherhood", maybe both groups are similar in that both like spending time with their children. Failing to take this into account would lead us to overestimate differences between groups. Once again, singular value decomposition is able to address this issue. In the example we just mentioned, we indeed find that one of the main dimensions loads heavily on both "fatherhood" and "motherhood", implying that both interests are reflective of the same underlying preference dimension.

Another concern is how Facebook deals with different languages. Each interest has a numerical identifier, so that users who are interested in "pain" in France, "pan" in Spain, "хлев" in Russia, etc., will all get identified as being interested in "bread". Of course, this does not erase all differences between languages. For example, one language may have separate concepts for "fatherhood" and "motherhood", whereas another only has a common concept for "parenthood". Singular value decomposition is likely to resolve this issue too. In this particular example, we find that the dimension that loads heavily on "fatherhood" and "motherhood" also loads heavily on "parent" and "parenting", indicating that these interests are associated with the same underlying preference.

[8] A notable country absent from our sample is China, where Facebook penetration is less than 1%.

Fig. 3. Gender equality and difference in interests between men and women.
Figure depicts the 2018 Gender Gap Index of the World Economic Forum (WEF) on the horizontal axis and the cosine distance between the vectors of 45,397 Facebook interest frequencies of men and women on the vertical axis for the baseline sample of countries (population > 1 million and Facebook penetration > 0.25).

3. Differences in interests between men and women

This section takes a first look at the data by exploring the cross-country relation between gender equality and the gender gap in all interests.

Raw correlation. For countries with a population above one million and a Facebook penetration rate of at least 25%, we have Facebook data on 106 countries. Of those countries, 98 also have data on gender equality. Focusing on this baseline sample, Fig. 3 depicts the raw correlation between gender equality and the difference in interests between men and women, computed as the cosine distance between all 45,397 interests. The correlation is slightly positive, suggesting that men and women in more gender-equal countries exhibit slightly larger differences in interests.

Partial correlation. To control for confounding determinants, we take a regression approach. Our baseline estimating equation is

$$CosDist_c = \beta GenderEq_c + \gamma Z_c + \varepsilon_c \tag{2}$$

where $CosDist_c$ is the cosine distance between the vectors of 45,397 interest frequencies of men and women in country i, $GenderEq_c$ is gender equality, Z_c is a vector of controls, and ε_c is an error term. Our main coefficient of interest is β, the partial correlation between gender equality and the difference in interests between women and men.

Table 1 reports the results for seven different specifications. Column (1) is our baseline specification: in addition to gender equality, it includes GDP per capita and the entropy of interests as regressors. As in Fig. 3, we find a weak positive relation between gender equality and the difference in interests between men and women. However, the corresponding coefficient is not statistically significant at the 10% level. The other control variables show that economic development is associated with larger gender differences, in line with Falk and Hermle (2018), whereas greater diversity in interests is associated with smaller gender differences.

In columns (2) through (5), we add further controls: regional dummies, religious composition, Soviet influence, geo-climatic conditions, population size, and Facebook penetration. We do not find any strong relation between gender equality and preference differences between men and women. The last two regressions in columns (6) and (7) return to the most basic specification, but use alternative measures of gender equality from the OECD and the UNDP. Although there we find a statistically significant effect of gender equality on the gender interest gap, our findings are overall inconclusive.

4. Difference between gendered and non-gendered interests

In this section, we start by classifying the 45,397 interests into two groups, those that are systematically related to gender and those that are not. We then analyze the relation between gender equality and gender differences for each of these two groups of interests.

Table 1
Gender differences in interests and gender equality: All interests.

Dependent variable: Cosine distance between men and women based on 45,397 Facebook interests							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Gender equality (WEF)	0.050	−0.017	0.090*	−0.020	0.021		
	(0.035)	(0.042)	(0.051)	(0.045)	(0.042)		
Gender equality (OECD)						0.001***	
						(0.000)	
Gender equality (UNDP)							0.048**
							(0.019)
Log GDP per capita	0.008***	0.006***	0.008***	0.005*	0.010***	0.006**	0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
Entropy	−0.059***	−0.060***	−0.054***	−0.047**	−0.049**	−0.060***	−0.055***
	(0.021)	(0.017)	(0.016)	(0.019)	(0.021)	(0.022)	(0.020)
WB Regional Dummies		✓					
Religion & Ideology			✓				
Geography & Climate				✓			
FB Penetration & Size					✓		
Observations	98	98	95	91	98	82	98
R^2	0.234	0.494	0.523	0.475	0.286	0.335	0.281

Robust standard errors in parentheses. The dependent variable is the difference between men and women based on 45,397 Facebook interests. The sample consists of countries with population > 1 million and Facebook penetration > 0.25. All regressions include a constant. WB Regional Dummies refer to dummy variables for sub-Saharan Africa, Middle East and North Africa, Europe and Central Asia, East Asia and the Pacific, North America, and Latin American and the Caribbean; Religion & Ideology refer to share of Protestants, share of Catholics, share of Muslims, and a dummy variable for Soviet influence; Geography & Climate refer to area (logarithm), land suitability, terrain roughness, temperature, and precipitation; FB Penetration & Size refer to population (logarithm) and the share of Facebook penetration.

* $p < 0.10$.

** $p < 0.05$.

*** $p < 0.01$.

4.1. Gendered and non-gendered interests

Starting with all 45,397 interests, we define two subsets of interests, one that is gendered and one that is non-gendered. We call an interest gendered if in more than 90% of countries the interest is more frequent among one of the genders. Examples include “engineering”, “fatherhood”, “romantic comedies”, “hunting” and “baking”. We call an interest non-gendered if in at least 30% of countries the interest is more frequent among men and in at least another 30% of countries it is more frequent among women.[9] Examples include “language school”, “blood donation” and “positive attitude quotes”. This procedure yields 2685 gendered interests and 2920 non-gendered interests.[10] For each country, we compute two cosine distances between men and women, one based on the set of gendered interests and another based on the set of non-gendered interests. Online Appendix Table C.2 provides these distance measures for all countries in our sample. When mapping these distances in Fig. 4, we notice that many countries where men and women are relatively similar in non-gendered interests show relatively large differences in gendered interests.

Before empirically analyzing how gender equality impacts the gap between men and women for these two types of interests, it is important to assess whether our classification is reasonable beyond simply giving some examples. To that end, we examine whether our list of gendered interests aligns with expected gender stereotypes. For this, we turn to large language models (LLMs) which are well-equipped for natural language processing tasks such as classification. Of the 2685 gendered interests, approximately 1500 show a worldwide male bias, and approximately 1000 exhibit a worldwide female bias. For each of these two groups, we prompt ChatGPT to classify the interests using a few keywords. For the male-biased interests, it identifies six main themes: “automotive and transportation”; “technology and gadgets”; “sports and physical activities”; “entertainment and media”; “gaming and e-sports”; and “business and finance”. And for the female-biased interests, it returns: “fashion and beauty”; “parenting and family”; “food and cooking”; “health and wellness”; “entertainment and celebrities”; and “home and lifestyle”.[11]

Reassuringly, this categorization of male- and female-biased interests matches standard gender stereotypes. More specifically, the male-biased interests reflect well-documented patterns: men’s tendency to gravitate toward technical fields (Stoet and Geary, 2018; Aldén and Neuman, 2022), a greater propensity for competition (Croson and Gneezy, 2009; Niederle and Vesterlund, 2011; Samek,

[9] To have a sufficiently broad cross-section of countries, we apply the procedure to the sample of 131 countries with a population > 1 million, Facebook penetration > 2.5% and Facebook users > 100,000, for which we have data on gender equality from the World Economic Forum.

[10] This leaves many unclassified interests, which fall into two main groups. One group includes interests that are weakly gender-biased but do not reach the 90% threshold. The other group consists of interests with zero audience among both men and women in a large number of countries. Many of these are local interests relevant only in one or a few countries.

[11] For each one of these themes, ChatGPT also returns illustrative examples. Online Appendix Table C.3 provides further details.

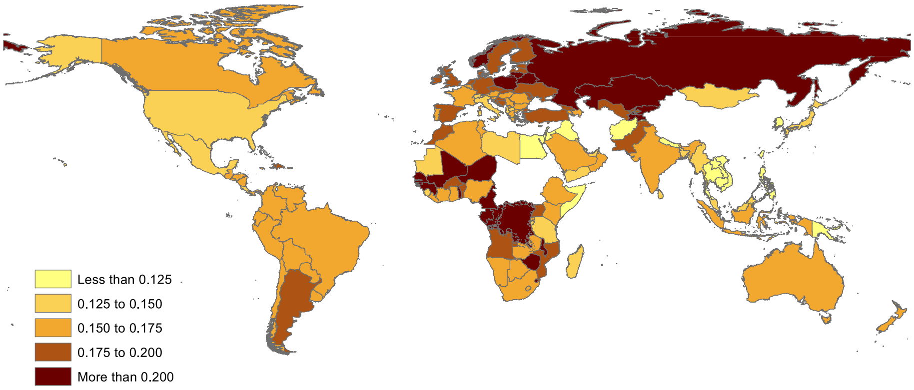

(a) Gendered Interests

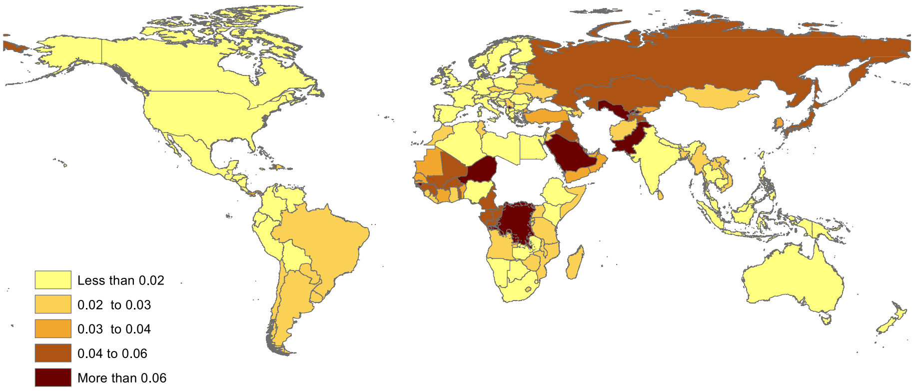

(b) Non-Gendered Interests

Fig. 4. Differences between men and women in gendered and non-gendered interests.
Figure shows the cosine distance between the interest frequency vector of men and women based on interests that are more frequent in one of the genders in at least 90% of countries (Panel a) and based on interests that are more frequent in men in at least 30% of countries and more frequent in women in at least 30% of countries (Panel b). Sample is restricted to countries with population above one million, Facebook penetration rate above 2.5%, and number of Facebook users greater than 100,000.

2019; Holmlund et al., 2023), their overrepresentation in finance (Bertrand, 2011), and greater emphasis on physical strength (Wood and Eagly, 2012). In contrast, the female-biased interests are consistent with the persistent sex-based division of labor, whereby women do most of core household and childcare tasks (Charles, 2011; Wood and Eagly, 2012), as well as with women's greater inclination to work with people rather than things (Su et al., 2009; Kuhn and Wolter, 2022).

As for the list of non-gendered interests, it is less clear what they should capture. When prompted with the same question for this list, ChatGPT identifies five main themes: "culture and arts"; "travel and geography'; "religion and spirituality"; "social and political topics"; and "science and technology". While this last category may have some overlap with the male-biased category "technology and gadgets", it is broader, encompassing topics that do not exhibit a male bias, such as "biology" and "sustainability".

The gender gap in gendered vs non-gendered interests. Fig. 5 depicts the raw correlation between gender equality and differences between men and women in gendered and non-gendered interests. The difference is immediately apparent: greater gender equality is associated with larger differences in gendered interests (Panel (a)), but smaller differences in non-gendered interests (Panel (b)).

Table 2 does a more in-depth analysis of these relations, based on the same seven regressions as Table 1. When comparing our findings for gendered interests in Panel A to those for non-gendered interests in Panel B, we observe the same stark difference as in the scatter plots. For gendered interests, there tends to be a strong positive association between gender equality and differences in interests between men and women, whereas for non-gendered interests, there tends to be a strong negative association between the two. That is, gendered interests diverge with gender equality, whereas non-gendered interests converge. The magnitudes of the effects are large: in the most basic specification in column (1), a one standard deviation increase in gender equality is associated with a 35% standard deviation increase in the difference between men and women in gendered interests, and a 46% standard deviation decrease in their difference in non-gendered interests.

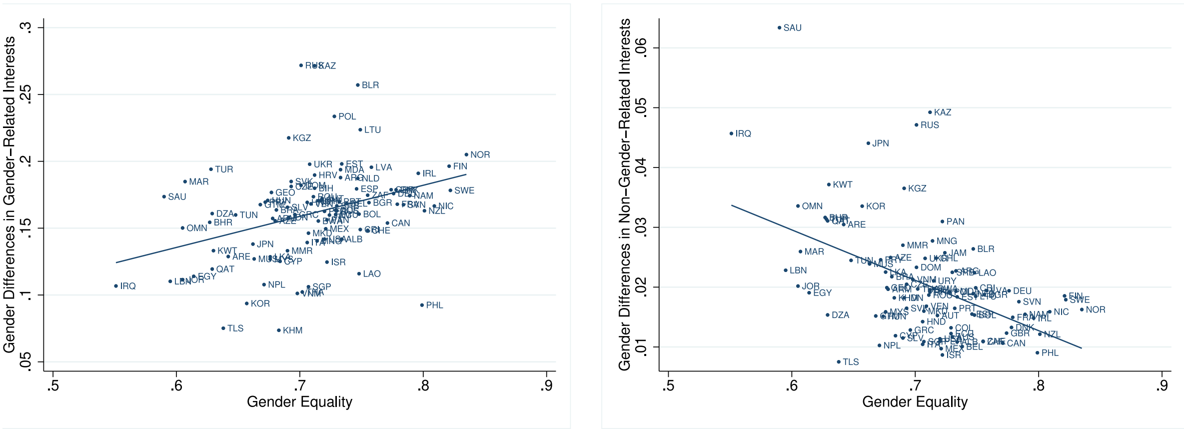

(a) Gendered Interests (b) Non-Gendered Interests

Fig. 5. Gender equality and differences in gendered vs. non-gendered interests.
Figure depicts scatter plots of gender equality on the horizontal axis and differences in interests between men and women, using the baseline sample of countries (population > 1 million and Facebook penetration > 0.25). Panel (a) shows the differences between men and women in gendered interests (i.e., interests that are more frequent in one of the genders in at least 90% of countries), whereas Panel (b) shows the differences between men and women in non-gendered interests (i.e, interests that are more frequent in men in at least 30% of countries and more frequent in women in at least 30% of countries).

Online Appendix Table C.4 considers the same seven specifications when expanding the sample to include countries with a population below one million. The results become slightly stronger. Online Appendix Table C.5 does the same, but for a sample that lowers the Facebook penetration threshold from 25% to 5%. The results are slightly weaker, but do not change qualitatively. In Table 2, the level of development tends to be associated with a larger gender gap, especially for non-gendered interests. When using the human development index (HDI) as an alternative measure of development, results are similar, as shown in Online Appendix Table C.6.

4.2. Robustness to classification of interests

Different thresholds. For an interest to be classified as gendered, we required it to have a common gender bias in at least 90% of countries. When rerunning the specification in column (1) of Table 2 for 25 different thresholds between 70% to 95%, the effect of gender equality is always positive and statistically significant at the 1% level. For the case of non-gendered interests, we required the interest to be more frequent among men in at least 30% and more frequent among women in at least another 30% of countries. When varying the threshold from 15% to 45%, the effect of gender equality is always negative and statistically significant at the 1% level. From this we conclude that our results are robust to less and more strict ways of classifying gendered and non-gendered interests.

Specific interests. As already mentioned, the gender-equality paradox has been linked to certain specific interests, such as the tendency to study STEM fields. Some other interest dimensions, such as family and relationships, also show a significant gender bias. As a way of validating our use of Facebook data, we assess whether the gender-equality paradox also shows up when restricting our attention to some of these specific gendered interests.

We start by examining the interest in STEM fields, focusing on the top-three types of undergraduate engineering degrees in the U.S. (“mechanical engineering”, “electrical engineering” and “civil engineering”),[12] along with “computer science” and “mathematics”. For each of these Facebook interests, we define the dependent variable as the difference in the interest share of men and women, and we re-run our basic specification, corresponding to column (1) in Table 2. As reported in Online Appendix Table C.7 Panel A, we find evidence of the gender-equality paradox for four out of the five interests. Only for “mathematics” do we not find a statistically significant relationship between gender equality and the interest gap between men and women. Overall, these results are consistent with the literature. While the gender-equality paradox is often associated with STEM in general, it is more commonly observed in engineering and computer science than in mathematics. Guiso et al. (2008), for instance, find that the gender gap in math is smaller in more gender-equal countries, although Anghel et al. (2020) note that this result depends on a country’s level of development. Similarly, Cheryan et al. (2017) document that gender differences are large in engineering and computer science (fields where women make up less than 20% of undergraduate degrees) but negligible in mathematics (where women make up more than 50% of undergraduate degrees).

Next, we explore Facebook interests related to family and relationships. Among interests with a Facebook audience of at least 100 million worldwide, we focus on some of the more popular ones related to family and relationships: “parenting”, “marriage”,

[12] See https://www.nsf.gov/nsb/sei/edTool/data/engineering-01.html.

Table 2
Differences between men and women in gendered vs non-gendered interests.

Panel A: Cosine distance between men and women based on gendered interests							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Gender equality (WEF)	0.232***	0.126**	0.203**	0.137*	0.168**		
	(0.061)	(0.062)	(0.097)	(0.079)	(0.072)		
Gender equality (OECD)						0.002***	
						(0.000)	
Gender equality (UNDP)							0.086**
							(0.036)
Log GDP per capita	0.006	0.003	0.006*	0.000	0.010**	0.003	−0.001
	(0.004)	(0.003)	(0.003)	(0.004)	(0.005)	(0.005)	(0.005)
Entropy	−0.048	−0.065**	−0.047*	−0.026	−0.028	−0.046	−0.037
	(0.034)	(0.026)	(0.026)	(0.031)	(0.034)	(0.037)	(0.034)
Regional Dummies		✓					
Religion & Ideology			✓				
Geography & Climate				✓			
FB Penetration & Size					✓		
Observations	98	98	95	91	98	82	98
R^2	0.180	0.563	0.518	0.397	0.249	0.271	0.115
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel B: Cosine distance between men and women based on non-gendered interests							
Gender equality (WEF)	−0.080***	−0.069***	−0.057***	−0.097***	−0.074***		
	(0.017)	(0.019)	(0.018)	(0.016)	(0.018)		
Gender Equality (OECD)						−0.000**	
						(0.000)	
Gender Equality (UNDP)							−0.029***
							(0.009)
Log GDP per Capita	0.003***	0.003***	0.003***	0.003***	0.003***	0.002***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Entropy	−0.031***	−0.035***	−0.036***	−0.032***	−0.034***	−0.033***	−0.038***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)
Regional Dummies		✓					
Religion & Ideology			✓				
Geography & Climate				✓			
FB Penetration & Size					✓		
Observations	98	98	95	91	98	82	98
R^2	0.529	0.567	0.621	0.637	0.544	0.518	0.428

Robust standard errors in parentheses. The dependent variable is the difference between men and women based on either the subset of interests that are more frequent in one of the genders in at least 90% of countries (Panel A) or on the subset of interests that are more frequent in men in at least 30% of countries and more frequent in women in at least 30% of countries (Panel B). The sample consists of countries with population > 1 million and Facebook penetration > 0.25. All regressions include a constant. WB Regional Dummies refer to dummy variables for sub-Saharan Africa, Middle East and North Africa, Europe and Central Asia, East Asia and the Pacific, North America, and Latin American and the Caribbean; Religion & Ideology refer to share of Protestants, share of Catholics, share of Muslims, and a dummy variable for Soviet influence; Geography & Climate refer to area (logarithm), land suitability, terrain roughness, temperature, and precipitation; FB Penetration & Size refer to population (logarithm) and the share of Facebook penetration.

* $p < 0.10$.

** $p < 0.05$.

*** $p < 0.01$.

"child", "grandparent" and "friendship". We then assess how gender equality affects the preference gap between women and men. Since these interests generally exhibit a female bias, we define the dependent variable as the difference in the interest share between women and men. As seen in Online Appendix Table C.7 Panel B, the results are consistent with the gender-equality paradox: in more gender-equal countries there are greater preference differences between women and men when it comes to family and relationships. From this we conclude that one of our paper's main findings applies not only to the general class of gendered interests, but also to specific gendered interests that have been studied in the literature.

4.3. Robustness to age composition

Given that our analysis focuses on the subset of countries with a Facebook penetration rate of at least 25%, we are fairly confident that our data are broadly representative of the population groups that we are interested in. However, some biases may persist even when reaching relatively high levels of Facebook penetration. Probably, the one that should concern us most is the age bias, since social media users tend to be younger than the overall population. Accounting for this bias is important because age may be a determinant of the differences in interests between men and women. For example, if older men and women are more different than younger men and women, then the coefficient on gender equality would be biased upward if more gender-equal countries have a larger proportion of older Facebook users. Illustrating this possibility, Vishkin (2022) finds that the wider gender gap in chess participation in gender-equal countries is partly due to a greater weight of older generation in those countries.

One way to address this concern is to re-run the regressions of Table 2, controlling for the ratio of older to younger Facebook users. Based on Erikson and Erikson (1998), we take 40 years as the cutoff age between young and old. As can be seen in Online Appendix Table C.8, the results are unchanged. Controlling for the age ratio, in more gender-equal countries the difference between men and women is larger for gendered interests and smaller for non-gendered interests.

Another way to address this concern is to run separate regressions for the old and the young. This requires us to have interest frequency data by age group. Unfortunately, getting such data for all 45,397 interests would be extremely time-consuming, and goes beyond the scope of this paper. However, for 5000 randomly chosen interests, we obtained frequency data by country for both the old and the young. Using the same definitions as before, we identify which of these 5000 interests are gendered and which are not. We then compute four distance measures: the distance between old men and old women for gendered interests; the distance between young men and young women for gendered interests; and analogous measures for the old and the young applied to non-gendered interests. For each one of these distance measures, we run our standard set of regressions. Online Appendix Tables C.9 and C.10 report the results. Two findings stand out. First, the results for the old are almost identical to the results for the young, suggesting that age composition is not material to the paper's findings. Second, we confirm our central result for both the old and the young: as gender equality increases, men and women tend to diverge in gendered interests, and they tend to converge in non-gendered interests.

4.4. Causality

So far we have refrained from using causal language. A society's gender equality is potentially endogenous because of reverse causality: differences in preferences between genders may affect the degree of equality between men and women. It is not obvious in which direction this potential endogeneity would bias our coefficients. On the one hand, if men and women want different things from life, this might translate in less gender equality in certain outcomes. This would increase the coefficient on gender equality, hence strengthening our findings for gendered interests and weakening them for non-gendered interests. On the other hand, if men and women have different preferences, there may be more pressure for women's rights and female political empowerment, leading to greater gender equality. This would decrease the coefficient on gender equality, hence weakening our findings for gendered interests and strengthening them for non-gendered interests.

To address this potential endogeneity concern, we take two approaches. In a first approach. we use the earliest available version of our gender equality index. As such, in our baseline specification we replace the gender equality index of 2018 by the one of 2006. The idea is that there is less likely to be a reverse causality issue between today's differences in preferences and the gender equality index of almost 15 years ago.[13] Columns (2) and (5) in Table 3 report our findings for gendered interests and non-gendered interests. When comparing to the baseline regressions reported in columns (1) and (4), there is no significant difference in the coefficients on gender equality. This somewhat allays concerns about reverse causality. Needless to say, to the extent that the unobservable factors that led to the possible identification problem in the first place are correlated over time, reverse causality is still an issue.

In a second approach, we turn to instrument variable estimation. We use the year when women gained the right to vote as an instrument for today's degree of gender equality. Since constructing gender equality through the political process takes many years, the time elapsed since female suffrage is bound to be a good predictor of today's gender equality. In fact, the correlation between the two variables is -0.48 (with a p-value of 0.00). How long ago women gained the vote is of course likely to affect today's differences in preferences between men and women. We would expect this effect to be mediated by the degree of female political empowerment and acquired economic, social and economic rights and opportunities. Since all these mediating factors are captured by the gender equality index we use, the exclusion restriction is likely to be satisfied.

Columns (3) and (6) in Table 3 report our findings based on IV estimation. The coefficients on gender equality are slightly larger in absolute value terms when using IV than when using OLS. In addition, the F-statistics of the first stage are larger than the Stock-Yogo critical values for 10% maximal IV size, so we can reject the hypothesis that our instrument is weak. Overall, these findings suggest that we can give a causal interpretation to our main result: more gender equality leads to larger differences between men and women in gendered interests and smaller differences in non-gendered interests. However, we must be cautious with this interpretation, because this result is based on our baseline specification. When considering more comprehensive specifications, our IV strategy ceases to pass the weak instrument test.

[13] For a recent example of this approach, see Gächter and Schulz (2016).

Table 3
Gendered and non-gendered interests: Causality.

	Cosine Distance Men–Women					
	Gendered interests			Non-gendered interests		
	(1) OLS	(2) Lagged	(3) IV	(4) OLS	(5) Lagged	(6) IV
Gender equality (WEF)	0.232***		0.244*	−0.080***		−0.105***
	(0.061)		(0.139)	(0.017)		(0.028)
Gender equality (WEF, 2006)		0.202**			−0.072***	
		(0.081)			(0.026)	
Log GDP per capita	0.006	0.004	0.006*	0.003***	0.003***	0.003***
	(0.004)	(0.005)	(0.003)	(0.001)	(0.001)	(0.001)
Entropy	−0.048	−0.060*	−0.049**	−0.031***	−0.039***	−0.031***
	(0.034)	(0.035)	(0.021)	(0.005)	(0.006)	(0.004)
Constant	0.366	0.520*	0.362**	0.325***	0.383***	0.334***
	(0.277)	(0.305)	(0.173)	(0.043)	(0.050)	(0.035)
Observations	98	86	98	98	86	98
R^2	0.180	0.173	0.179	0.529	0.543	0.509
Cragg-Donald F			23.70			23.70
Stock-Yogo 10% max IV size			16.38			16.38

Robust standard errors in parentheses. The dependent variable is the difference between men and women based on either the subset of interests that are more frequent in one of the genders in at least 90% of countries (columns (1)–(3)) or on the subset of interests that are more frequent in men in at least 30% of countries and more frequent in women in at least 30% of countries (columns (1)–(3)). The sample consists of countries with population > 1 million and Facebook penetration > 0.25. Columns (1) and (4) are identical to column (1) in Table 1. Columns (2) and (5) use the gender equality index of the WEF of 2006. Columns (3) and (6) are based on IV regressions, using the year when women obtained the right to vote as instrument of the gender equality index of the WEF.

* $p < 0.10$.

** $p < 0.05$.

*** $p < 0.01$.

5. Latent gendered and non-gendered preference dimensions

When using thousands of interests to construct distance measures, “noise” may obscure the underlying structure of the data. One source of noise is “synonymy”, the possibility that different interests reflect the same underlying interests. For example, people interested in “cars” and people interested in “automobiles” should be classified as having common interests. Similarly, people interested in “motherhood” and others interested in “fatherhood” might have a common interest in their family. Failing to take this into account would tend to overestimate differences between populations. Another source of noise is “polysemy”, the possibility that the same interest has different meanings or connotations for different populations. For example, people interested in “trust” in a financial sense and people interested in “trust” in a confidence sense should be classified as having different interests. Failing to take this into account would tend to underestimate differences between populations. In addition to synonymy and polysemy, noise may also be introduced by people trying to avoid showing an interest in certain specific topics.

These problems are well known from the text classification and information retrieval literature. In that literature, each group is a text and each interest is a word, with each text being identified by its vector of word frequencies (Baeza-Yates and Ribeiro-Neto, 2011). Retrieval techniques that match queries to documents need to compute distances between documents. The conventional methodology to deal with these issues is latent semantic indexing (LSI). It uses singular value decomposition (SVD), a method similar to principal component analysis, to create a lower-dimensional semantic space that places words that occur in similar documents close to one another (Deerwester et al., 1990).[14]

Applied to our problem, we use SVD to construct a lower-dimensional space that classifies interests held by populations with similar underlying interests as being closely related. Doing so allows us to get rid of redundant data, to focus on the main associative patterns in the Facebook interest data, and to address the problems of synonymy and polysemy. Before providing further details, it may be useful to provide some intuition for how this methodology is able to deal with synonymy and polysemy. For the case of synonymy, consider two groups, one with a strong interest in “cars” and another with a strong interest in “automobiles”. In higher-dimensional space, these two groups will appear less similar than they really are. In contrast, SVD recognizes that “car” and “automobile” belong to the same underlying concept, so both groups will be mapped closely together in lower-dimensional space. For the case of polysemy, for one group the interest “trust” may co-occur with other interests such as “finance”, “savings” or “investment”, whereas for another group the interest “trust” may co-occur with “confidence”, “faith” or “reliance”. With SVD,

[14] Singular value decomposition maximizes the value of the second moment of the projections of the uncentered data, whereas principal component analysis maximizes the variance of the projected data. In our case, the two methods produce very similar results.

the different meanings of the polysemous interest "trust" are likely to get separated into different latent dimensions, ensuring that both groups appear as dissimilar.

In the remainder of this section, we use two SVD-based approaches to identify the main gendered and non-gendered preference dimensions, and to examine the presence of the gender-equality paradox for these different latent dimensions. In the first approach, we apply SVD separately to the sets of gendered and non-gendered interests. We then compute the cosine distance between men and women based on their positions along the main gendered and non-gendered dimensions. Using these distances, we revisit the relationship between gender equality and preference differences between men and women. In the second approach, we apply SVD to the entire set of interests, without distinguishing between gendered and non-gendered interests. We then compare the positions of men and women across countries along the main latent dimensions. If, for a given dimension, men and women form distinct clusters, we consider it a gendered dimension; otherwise, we consider it non-gendered. Using this distinction, we again revisit how differences between men and women relate to the level of gender equality.

5.1. Dimensionality reduction of gendered and non-gendered interests

In this subsection we apply SVD separately to the sets of gendered and non-gendered interests. We then assess whether our main findings hold when focusing on these latent preference dimensions, rather than on the interests themselves.

Singular value decomposition. Consider the $I \times G$ interests-by-group matrix X, where the rows correspond to the I interests and the columns to the G country-gender groups. Not to overburden notation, I corresponds alternatively to gendered interests or to non-gendered interests, depending on the set of interests we are applying SVD to. Element x_{ig} of the matrix refers to the share of interest i in group g.[15] We denote the rank of matrix X by r, where $r \leq G$. A well-known theorem of linear algebra says that X can be decomposed as

$$X = U \Sigma V^T \tag{3}$$

where U is an orthogonal $I \times G$ matrix, Σ is an $G \times G$ diagonal matrix, and V^T is an orthogonal $G \times G$ matrix.[16] The first r diagonal elements of Σ correspond to the square roots of the r non-zero eigenvalues of XX^T. They are referred to as the non-zero singular values and they are ordered such that $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r$. The first r columns of U contain the orthonormal eigenvectors corresponding to the non-zero eigenvalues of XX^T. They are referred to as the left singular vectors. The first r columns of V contain the orthonormal eigenvectors corresponding to the non-zero eigenvalues of X^TX. They are referred to as the right singular vectors.

The goal of SVD is to discover the main latent or underlying interest dimensions. It may be useful to provide some intuition of how these dimensions are related to the matrix decomposition in (3). The columns of the $I \times G$ matrix U relate the different Facebook interests to each one of the latent interest dimensions. For example, the elements of the first column of U give the relative weights of each Facebook interest in the first dimension. The diagonal elements of the $G \times G$ matrix Σ then give a measure of the importance of each interest dimension. As they are declining in order, the first dimension is more important than the second, and so on. The columns of the $G \times G$ matrix V relate the different country-gender groups to each one of the latent interest dimensions. For example, the elements of the first column of V give the importance that each country-gender group attaches to the first interest dimension.

Dimensionality reduction. When computing distances between populations, it is useful to consider a reduced set of latent dimensions, rather than the full dimensionality of interests. By doing so, we get rid of noise and focus on the main associative patterns in the Facebook interest data.

To reduce the dimensionality of X to $\hat{r} < r$, we keep the first $\hat{r}$ singular values in Σ and their corresponding singular vectors in U and V. This yields

$$X_{\hat{r}} = U_{\hat{r}} \Sigma_{\hat{r}} V_{\hat{r}}^T \tag{4}$$

where $X_{\hat{r}}$ is an $I \times G$ matrix, $U_{\hat{r}}$ is an $I \times \hat{r}$ matrix, $\Sigma_{\hat{r}}$ is an $\hat{r} \times \hat{r}$ matrix, and $V_{\hat{r}}^T$ is an $\hat{r} \times G$ matrix. The matrix $X_{\hat{r}}$ is the best $\hat{r}$-rank approximation of X in the sense that it minimizes the sum of squared errors (Eckart and Young, 1936).

To find a value for $\hat{r}$, we plot the singular values in decreasing order, and keep all singular values before there is a large drop in the plot. This ad-hoc approach is referred to as identifying an "elbow" in the curve of singular values. Applying this method to the subset of gendered interests and on the subset of non-gendered interests yields a rank $\hat{r} = 8$ for both matrices. We refer to the truncated gendered and non-gendered matrices as, respectively, X_8^g and X_8^{ng}.

Revisiting the gender-equality paradox in a lower-dimensional subspace. For each of these two matrices, we recompute the cosine distances between men and women and re-run the same regressions as before. Table 4 presents the results, with Panel A corresponding to the gendered preference dimensions and Panel B to the non-gendered preference dimensions. We observe the same stark difference: more gender-equal societies exhibit larger differences between men and women along gendered dimensions and smaller differences along non-gendered dimensions. The magnitudes of the effects are similar to the ones we found before: in the most basic specification in column (1), the standardized β corresponding to gender equality is 40% in the case of gendered interests, and -43% in the case of non-gendered interests.

[15] More specifically, x_{ig} is defined as the share of signals expressed by group g that corresponds to interest i, i.e., $f_{gi}/\sum_i f_{gi}$, where f_{gi} is the number of users of group g that hold interest i. An alternative would be to define x_{ig} as the share of users in group g who are interested in i. We prefer the former measure because the number of interests per capita often differs substantially between genders within the same country.

[16] For an exposition, see, for example, Shores (2007).

Table 4
Gender differences based on main latent gendered and non-gendered dimensions.

Panel A: Cosine distance between men and women based on SVD of gendered interests							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Gender equality (WEF)	0.252***	0.125**	0.178*	0.170**	0.180***		
	(0.057)	(0.054)	(0.090)	(0.075)	(0.066)		
Gender equality (OECD)						0.002***	
						(0.000)	
Gender equality (UNDP)							0.110***
							(0.035)
Log GDP per capita	0.006*	0.003	0.006**	0.000	0.010**	0.002	−0.003
	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.005)
Entropy	−0.006	−0.016	0.001	0.017	0.019	−0.001	0.010
	(0.027)	(0.017)	(0.019)	(0.026)	(0.027)	(0.030)	(0.028)
Regional Dummies		✓					
Religion & Ideology			✓				
Geography & Climate				✓			
FB Penetration & Size					✓		
Observations	98	98	95	91	98	81	97
R^2	0.236	0.677	0.588	0.439	0.320	0.338	0.162
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel B: Cosine distance between men and women based on SVD of non-gendered interests							
Gender equality (WEF)	−0.022***	−0.018**	−0.018***	−0.014**	−0.018**		
	(0.008)	(0.007)	(0.006)	(0.005)	(0.007)		
Gender equality (OECD)						−0.000*	
						(0.000)	
Gender equality (UNDP)							−0.012***
							(0.003)
Log GDP per capita	0.001***	0.001***	0.001***	0.001***	0.001**	0.000*	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Entropy	−0.003***	−0.005***	−0.004***	−0.006***	−0.005***	−0.002**	−0.005***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.002)
Regional Dummies		✓					
Religion & Ideology			✓				
Geography & Climate				✓			
FB Penetration & Size					✓		
Observations	98	95	91	98	98	81	97
R^2	0.238	0.288	0.360	0.390	0.287	0.082	0.209

Robust standard errors in parentheses. The dependent variable is the difference between men and women based on the first eight dimensions of SVD on either subset of interests that are more frequent in one of the genders in at least 90% of countries (Panel A) or subset of interests that are more frequent in men in at least 30% of countries and more frequent in women in at least 30% of countries (Panel B). The sample consists of countries with population > 1 million and Facebook penetration > 0.25. The seven specifications are identical to those in Table 1. All regressions include a constant. WB Regional Dummies refer to dummy variables for sub-Saharan Africa, Middle East and North Africa, Europe and Central Asia, East Asia and the Pacific, North America, and Latin American and the Caribbean; Religion & Ideology refer to share of Protestants, share of Catholics, share of Muslims, and a dummy variable for Soviet influence; Geography & Climate refer to area (logarithm), land suitability, terrain roughness, temperature, and precipitation; FB Penetration & Size refer to population (logarithm) and the share of Facebook penetration.

* $p < 0.10$.
** $p < 0.05$.
*** $p < 0.01$.

Rather than considering the cosine distance between men and women based on the sets of gendered and non-gendered dimensions, we could also examine each dimension separately. That is, we could calculate sixteen distances between men and women—one for each of the eight gendered and eight non-gendered preference dimensions. When rerunning our basic specification separately for each gendered dimension, the coefficient on gender equality has the expected positive sign for six out of the eight dimensions; and for all eight non-gendered dimensions, the coefficient has the expected negative sign. Thus, our key findings tend to hold dimension by dimension. Furthermore, the interpretation of each dimension generally aligns with our understanding of gendered and non-gendered preferences. For instance, ChatGPT characterizes the interests along one of the gendered dimensions as ranging from “functional, scientific, or educational themes” to “social identity, feminized roles, and aesthetic or emotional cues”, and it characterizes the interests along another gendered dimension as ranging from “nature, leisure, aesthetic, and gentler cultural references” to “conflict-oriented, dramatic, or media-dense content”.

5.2. Identifying gendered and non-gendered preference dimensions

In this section, we propose an alternative SVD-based approach for distinguishing between gendered and non-gendered preferences. Instead of first classifying interests into gendered or non-gendered categories and then applying dimensionality reduction, we use SVD on all 45,397 interests and classify the resulting preference dimensions as gendered or non-gendered, based on whether men and women form distinct clusters or not. We then analyze whether there are systematic differences in the relation between gender equality and gender differences in preferences depending on whether the preference dimension is gendered or not.

Detecting gendered and non-gendered preference dimensions through svd. Along each of the latent preference dimensions identified by singular value decomposition of the data matrix containing all 45,397 interests, we can position men and women of different countries. Starting off with the $I \times G$ interest by country-gender matrix X of rank r, SVD gives us $X = U\Sigma V^T$. The matrix ΣV^T places country-gender groups in the vector space of rank r. More specifically, the non-zero first r rows of the $G \times G$ matrix ΣV^T give the positions of the country-gender groups along each one of the r interest dimensions. For example, the elements of the first row give the positions of men and women in different countries along the first preference dimension. The position of country-gender group g along preference dimension i can be written as $\sigma_i v^T_{ig}$, where v^T_{ig} is the element corresponding to row i and column g of matrix V^T. Applying the "elbow" method, we identify the first nine dimensions as being the most relevant.

To visualize the relative positions of men and women in the different countries, Fig. 6 displays two-dimensional scatter plots for each one of the first nine preference dimensions, with the position of women on the horizontal axis and the position of men on the vertical axis. Consider, for example, the scatter plot that depicts the preference dimension associated with the second singular vector V_2. Each point corresponds to one country, and gives the position of women in that country on the horizontal axis and the position of men in that country on the vertical axis. Points that are above the 45° line refer to countries where the position of men along preference dimension 2 is higher than that of women.

To distinguish between gendered preference dimensions and non-gendered preference dimensions, we start with a visual inspection of the different panels of Fig. 6. Of the different dimensions, the one associated with singular vector V_4 displays the strongest gender component: independently of country, women have a positive value while men have a negative value. Along that dimension women of different countries tend to be more similar to each other than to men of their own country. In other words, women and men form distinct clusters. To further illustrate how V_4 captures a dimension along which men's and women's interests are very different, we can multiply the fourth left singular vector U_4 by the fourth singular value σ_4 to obtain the position of each one of the 45,397 different interests along the fourth preference dimension. The interests with lower values correspond to "masculine" interests, and the ones with higher values to "feminine" values. Among the most masculine interests, many relate to cars, technology, gaming and sports, and among the feminine interests, many relate to cooking, family and fashion.[17] The dimensions associated with singular vectors V_2 and V_5 also display a slight gender bias: men either have systematically higher values than women (V_2), or the other way around (V_5). However, in contrast to V_4, along V_2 and V_5 men and women do not form distinct clusters. Instead, along these dimensions, there continues to be an important country component: women of a particular country tend to be closer to men of their own country than to women in other countries, although in each country men and women are systematically different. The other dimensions V_3, V_6, V_7 and V_8 do not show a clear gender bias, and can be considered to be mostly unrelated to gender. For these dimensions, some points are above and others are below the 45° line.[18] While our main focus is on gender, it is worth considering what these non-gendered dimensions represent. Although interpreting these dimensions is not straightforward, some capture specific characteristics of the world's main cultural groups. For instance, in dimension V_3 all Slavic countries cluster at one extreme, implying that this dimension likely reflects certain features or traits specific to Slavic culture.

To more formally identify gendered preferences and non-gendered preferences, for each dimension we compute the incremental R^2 from adding gender to a regression of the positions of men and women on a full set of country dummies. The greater the explanatory power of gender, the larger the incremental R^2. This methodology confirms our visual inspection of Fig. 6. Of the different dimensions, the incremental R^2 due to gender for V_4 is 88%. For V_2 and V_5, the incremental R^2 is between 2% and 4%, and for all other dimensions it is below 1%. We can therefore conclude that dimension V_4 is gendered, dimensions V_2 and V_5 are weakly gendered, and dimensions V_3, V_6, V_7, V_8 and V_9 are non-gendered.

Gender differences along gendered and non-gendered preference dimensions. Next, we analyze whether there is a difference in the relation between gender equality and gender differences in preferences depending on whether the preference dimension is gendered or not. We compute the Euclidean distance between men and women based on gendered and non-gendered dimensions, and re-run

[17] For a full list of the 500 most masculine and the 500 most feminine interests along preference dimension 4, see Online Appendix Tables C.11 and C.12. Examples of the most masculine interests include Automobiles, BMW, Motorcycles, Personal finance, War, Vladimir Putin, Game Consoles, Free Software, Engine, SUVs, Cameras, Outdoor recreation, UEFA Champions League, Lionel Messi, Sport cars, Wheel, Bluetooth, Martial arts, Hunting, Military, Tool, Poker, Shooter games, Computer monitors. Examples of the most feminine interests include Dresses, Cosmetics, Infant, Motherhood, Poetry, Beauty salons, Pregnancy, Boutiques, Child, Cooking, Cake, Chocolate, Jewelry, Handbags, Blouse, Hairstyle, Weddings, Recipes, Make-up artist, Skirt, Cuisine, Skin, Flower, Childbirth, Wedding dress, Weight loss, Psychology, Yoga, Breastfeeding.

[18] In our description of the different dimensions, we did not mention V_1. Along that dimension, all countries and genders present very similar values. This is the dimension that captures the mean positions. It can also be interpreted as the dimension that captures the preferences common to all groups. In other words, it is the dimension that measures the common humanity of all groups, and it therefore is not useful as a way of separating groups (countries and/or genders). In principal component analysis, this dimension is absent, because of data normalization. In the rest of the analysis, we will ignore V_1.

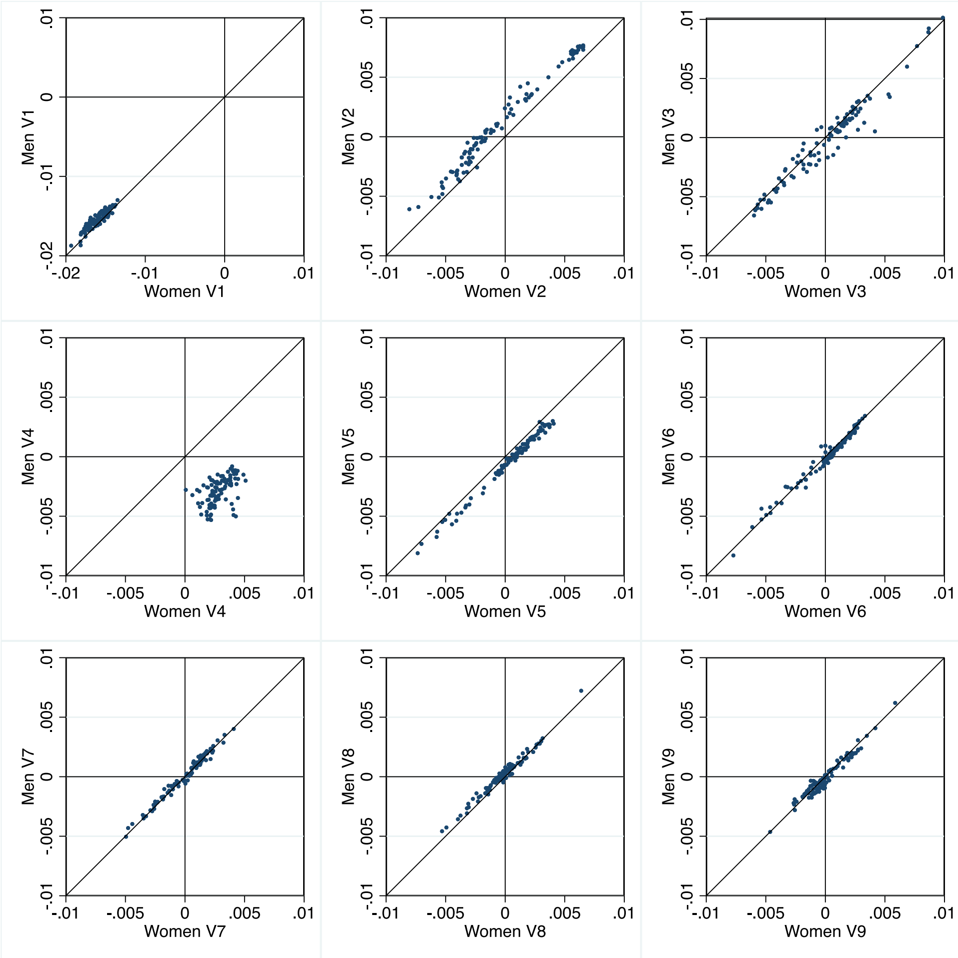

Fig. 6. Positions of women and men along main preference dimensions.
Figure shows the positions of women and men in the different countries along the nine most important preference dimensions as determined by SVD.

the same seven regressions as before.[19] Table 5 reports the results. Panels A and B focus on gendered dimensions (taking either a strict definition, based on V_4, or a more lenient definition, also including V_2 and V_5), whereas Panel C focuses on non-gendered dimensions. Once again, we confirm the paper's main finding. The coefficients on gender equality tend to be positive and statistically significant in Panels A and B, whereas they tend to be negative and statistically significant in Panel C. Hence, more gender-equal societies exhibit greater differences between men and women for gendered preferences and smaller differences between men and women for non-gendered preferences. This confirms the paper's main result: the gender-equality paradox applies to gendered preferences, but not to non-gendered preferences.

Specific interests. Because the notion of gendered preferences may be abstract, it is informative to examine whether the gender-equality paradox also applies to specific gendered interests. To avoid selecting specific interests in an ad-hoc manner, we rely on ChatGPT to generate a list of representative interests.

Specifically, we start by identifying the 5000 most male-biased interests and the 5000 most female-biased interests along V_4, and then prompt ChatGPT to classify these into broader categories using a few keywords. Online Appendix Table C.13 provides further details. For the male-biased interests, ChatGPT returns the categories "automotive and transportation", "technology and electronics", "sports and physical activities", "entertainment and media", "gaming and e-sports", and "lifestyle and fashion". For the female-biased interests, the categories are "fashion and beauty", "parenting and family", "food and cooking", "health and wellness", "entertainment and celebrities", and "home and lifestyle".

[19] We use the Euclidean distance, rather than the cosine distance, because in some cases the distance is based on just one dimension. In the cases for which the distance is based on more than one dimension, using the cosine distance yields qualitatively very similar results.

Table 5
Gendered and non-gendered dimensions based on SVD.

Panel A: Euclidean distance between men and women based on gendered dimension V_4							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Gender equality (WEF)	0.005***	0.002	0.005**	0.003	0.003*		
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)		
Gender equality (OECD)						0.000***	
						(0.000)	
Gender equality (UNDP)							0.004***
							(0.001)
Log GDP per capita	0.000***	0.000**	0.000***	0.000	0.000***	0.000	−0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Entropy	−0.003***	−0.003***	−0.003***	−0.003***	−0.002***	−0.003***	−0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Regional Dummies		✓					
Religion & Ideology			✓				
Geography & Climate				✓			
FB Penetration & Size					✓		
Observations	98	98	95	91	98	81	97
R^2	0.283	0.651	0.653	0.522	0.340	0.421	0.345
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel B: Euclidean distance between men and women based on gendered dimension V_2, V_4 and V_5							
Gender equality (WEF)	0.006***	0.002	0.006**	0.004	0.005**		
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)		
Gender equality (OECD)						0.000***	
						(0.000)	
Gender equality (UNDP)							0.004***
							(0.001)
Log GDP per capita	0.000***	0.000**	0.000***	0.000	0.000***	0.000	−0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Entropy	−0.003***	−0.003***	−0.003***	−0.003***	−0.003***	−0.003***	−0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Regional Dummies		✓					
Religion & Ideology			✓				
Geography & Climate				✓			
FB Penetration & Size					✓		
Observations	98	98	95	91	98	81	97
R^2	0.311	0.672	0.670	0.548	0.372	0.435	0.355
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel C: Euclidean distance between men and women based on non-gendered dimension V_3, V_6, V_7, V_8 and V_9							
Gender equality (WEF)	−0.005***	−0.003***	−0.004***	−0.004***	−0.004***		
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)		
Gender equality (OECD)						−0.000*	
						(0.000)	
Gender equality (UNDP)							−0.002***
							(0.000)
Log GDP per capita	0.000***	0.000***	0.000***	0.000***	0.000*	0.000	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Entropy	−0.001**	−0.001***	−0.001***	−0.001***	−0.001***	−0.001	−0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

(*continued on next page*)

Table 5 (*continued*).

Regional Dummies		✓					
Religion & Ideology			✓				
Geography & Climate				✓			
FB Penetration & Size					✓		
Observations	98	98	95	91	98	81	97
R^2	0.326	0.531	0.461	0.493	0.400	0.125	0.246

Robust standard errors in parentheses. The dependent variable is the Euclidean distance between men and women based on V_4 (Panel A), V_2, V_4 and V_5 (Panel B) and V_3, V_6, V_7, V_8 and V_9 (Panel C). The sample consists of countries with population > 1 million and Facebook penetration > 0.25. The seven specifications are identical to those in Table 1. All regressions include a constant. WB Regional Dummies refer to dummy variables for sub-Saharan Africa, Middle East and North Africa, Europe and Central Asia, East Asia and the Pacific, North America, and Latin American and the Caribbean; Religion & Ideology refer to share of Protestants, share of Catholics, share of Muslims, and a dummy variable for Soviet influence; Geography & Climate refer to area (logarithm), land suitability, terrain roughness, temperature, and precipitation; FB Penetration & Size refer to population (logarithm) and the share of Facebook penetration.

* $p < 0.10$.

** $p < 0.05$.

*** $p < 0.01$.

For each category, ChatGPT also provides examples of specific interests: for "automotive and transportation", it cites "BMW" and "motorsports"; for "sports and physical activities", it mentions "martial arts" and "racing"; for technology, it cites "smartphones" and "Xbox"; for "food and cooking", it mentions "French cuisine" and "desserts"; for "parenting and family", it cites "baby showers" and "preschool"; and for "health and wellness", it mentions "weight loss" and "yoga". Consistent with our discussion of the literature, these interests reflect the people vs. things divide, the male emphasis on physical strength and competition, and the gendered division of household and childcare responsibilities.

In all, we consider 18 male-biased and 18 female-biased specific interests, and re-run our basic specification for each. Online Appendix Table C.14 reports the findings. Of the 36 regressions, 30 yield the expected positive coefficient on gender equality, with 23 statistically significant at the 10% level. For example, the difference between men and women in their interest in football or preschool is larger in more gender-equal countries. Overall, these results confirm that the gender-equality paradox applies to specific gendered interests, in addition to the more abstract notion of gendered preferences.

6. Discussion and further analysis

In this section, we discuss how the paper's main findings relate to evolutionary psychology, social role theory, and gender essentialism. We also briefly explore possible implications for labor markets, intra-household bargaining, and political decision-making.

Evolutionary psychology, social role theory, and gender essentialism. Evolutionary psychology suggests men and women evolved differently to address distinct adaptive challenges (Atari et al., 2020). In societies with greater gender equality, men and women can express their innate preferences more freely, leading to wider preference differences (Buss, 1989; Schmitt, 2015). In contrast, social role theory posits that gender differences arise from socialization, norms, and power structures (Schmitt et al., 2017). As gender equality increases, one might expect these gender norms erode, loosening the constraints on how men and women should behave, thereby causing preference differences to narrow (Eagly et al., 2000). If so, these theories would seem to imply that in more gender-equal societies, differences between men and women should widen for innate preferences, while they should narrow for socially constructed preferences.

While there is no simple mapping of our preferences into innate and socially constructed, arguably innate preferences are more likely to display a systematic bias toward the same gender across the globe. If an interest, such as traveling, is more popular among men in some countries and more popular among women in other countries, it is unlikely to have an important innate component. But if an interest, such as sports or war, shows a universal male bias, it may have an innate basis. Using this probabilistic mapping, our findings align with this interpretation of evolutionary psychology and social role theory: in more gender-equal countries, differences between men and women are larger for gendered preferences (more likely to be innate) and smaller for non-gendered preferences (more likely to be socially constructed).

Though compelling, this interpretation comes with two important caveats. First, the fact that gendered interests display the same gender bias worldwide does not necessarily make them innate. Such interests may still be shaped by cultural or social forces. In fact, dual inheritance theory, which emphasizes the coevolution of genes and culture, suggests that disentangling the truly innate from the culturally determined may be impossible (Boyd and Richerson, 2005). The theory of gender essentialism further argues that cultural beliefs about innate gender differences can have the same impact as actual innate differences, reinforcing the difficulty of separating biology from culture (Charles and Bradley, 2009; England, 2010).[20] Therefore, while it may be tempting to interpret the systematic differences between men and women in gendered interests as being innate, in reality there is no easy way to discriminate between the roles of nature and nurture in shaping preferences.

[20] As expressed by sociologist Maria Charles in a 2020 interview with the GenderSci Lab at Harvard University: "In a world where pervasive beliefs about innate gender difference shape the lens through which people see and respond to the world around them, it is nearly impossible to accurately measure relative effects of biology [...] and sociocultural factors on occupational outcomes" (Richardson, 2020).

Second, gender essentialism also suggests that gender stereotypes may get reinforced in more gender-equal societies, meaning that social role theory need not imply a narrowing of preference differences. To understand this, it is useful to distinguish between two types of gender equality. Following the literature on the gender-equality paradox, our paper has focused on "vertical" gender equality—the belief associated with gender liberalism that men and women should have equal rights and opportunities. However, there is also the notion of "horizontal" gender equality—the belief associated with gender essentialism that men and women have inherently different characteristics, abilities, behaviors and preferences. It has been argued that greater vertical gender equality often coincides with lower horizontal gender equality (Charles, 2011; Cotter et al., 2011). As women and men achieve equality in rights and opportunities, they may fall back on gender stereotypes as a means of affirming their gender identity (Charles, 2011). This effect may be heightened in individualistic, post-materialist societies that prioritize self-realization and self-expression (Napp and Breda, 2022). As a result, social roles may become more – not less – gendered in societies with greater vertical gender equality, potentially explaining why such societies display larger gaps in gendered preferences.

Controlling for gender-essentialist norms and individualism. Given the argument that greater gender equality in rights and opportunities may reinforce gender essentialist beliefs, especially in individualistic societies, it is worth exploring this explanation in the context of our findings. Inspired by Breda et al. (2020), who find that the gender-equality paradox in math preferences disappears when controlling for gender-essentialist norms, we re-run our specifications of Table 2, adding controls for four different measures of gender essentialist norms. Three of these measures capture horizontal differences between men and women in ability, behavior and personality (Kosakowska-Berezecka et al., 2024), while the fourth measures the stereotype that "math is not for girls" (Breda et al., 2020). As reported in Online Appendix Table C.15, we do not find any evidence of gender essentialism explaining the larger gap between men and women in gendered interests. We also evaluate whether individualism might explain the gender-equality paradox. We do not find any evidence to that effect either (Online Appendix Table C.16).

Policy implications. In what follows, we briefly discuss some of the implications of our findings for gender segregation in the workplace, political decision-making, and intra-household bargaining.

If rising vertical gender equality is accompanied by greater horizontal gender equality, gender segregation in the workplace and wage disparities between men and women are likely to continue. This prompts the question of how policy should address the persistence of gender essentialist norms. The answer depends on the relative importance of nature and nurture in shaping these norms. As Levanon et al. (2016) argue, if gender differences stem from innate preferences, they may be justifiable on liberal-egalitarian grounds. In that case, even if job segregation results from gender essentialism, the wage penalty for women would be offset by the higher utility they derive from their work. However, if gender essentialist norms are socially constructed, policy interventions should aim to address and reduce their persistence.

Our findings also have implications for political representation and intra-household bargaining. As societies become more gender-equal, our results suggest that men and women will diverge more in gendered preferences. For instance, interests such as "parenting", "offspring", "kindergarten", "adolescence", "children's clothing" and "teacher" are strongly feminine worldwide, whereas interests such as "war", "police", "shooter games" and "army" are strongly masculine. These same differences are reflected in political decision-making. For example, Clots-Figueras (2012) finds that the election of women politicians in India improves educational attainment, while Lippmann (2021) shows that male legislative activity in the French parliament focuses more on the military. Similarly, Funk and Gathmann (2015) document that in Switzerland's direct democracy initiatives, women make different choices from men in environmental protection, defense spending and welfare policy. As societies become more gender-equal, our results show that women and men are likely to become more different in terms of these interests. As a result, when it comes to certain policies, such as those related to education and defense, it may be precisely in the most gender-equal countries where achieving gender balance in political representation is most critical. These gender differences also matter for the design of policies related to intra-household decisions. For example, Quisimbing and Maluccio (2020) demonstrate that giving more assets to women translates into an increase in spending on offspring.

7. Concluding remarks

This paper used information on the frequency of 45,397 Facebook interests to study how the difference in preferences between men and women changes with a country's degree of gender equality. The paper's main finding is that for interests or preferences that are gendered, we observe a larger gender gap in more gender-equal countries, whereas the opposite is true for interests or preferences that are non-gendered.

We established the paper's central finding by using many different ways of classifying interests and preferences. First, we split up all 45,397 interests into gendered and non-gendered interests, classifying as gendered the interests that are more frequent among the same gender for almost all countries. Second, we experimented with more stringent and more lenient thresholds when classifying interests as related to gender or not. Third, we used singular value decomposition on both subsets of interests to focus on the relevant latent dimensions. Fourth, we also used singular value decomposition on all interests, to then classify the resulting latent preference dimensions as related to gender or not. We found our paper's main result to be robust to these different ways of distinguishing between gendered and non-gendered interests and preferences.

While our results speak to competing theories about gender differences in preferences, such as evolutionary psychology and gender essentialism, they obviously do not resolve the debate on the relative importance of nature and nurture in shaping those differences. Our main contribution, instead, is to have established that gender differences in preferences display systematic patterns when considering thousands of interests. More specifically, we demonstrated that the gender-equality paradox in occupational and educational preferences is a more general phenomenon that extends to all gendered preferences. In contrast, for non-gendered interests the gender-equality paradox does not apply.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Interpreting Facebook interests as revealed preferences

By continuously tracking the online and offline activity of its Facebook users, Facebook is able to identify its users' interests. Before giving more details on this process, we provide a simple conceptual framework that helps us to understand the connection between the interests that Facebook assigns to users and user preferences.

An individual j gets utility from a set of I goods, services, actions, and beliefs, indexed by i:

$$u^j = u(x_1^j, x_2^j, \dots, x_I^j) \tag{5}$$

where x_i^j is the quantity of i enjoyed by j. For example, j might get utility from eating pizza, from having kids, from walks in the forest, or from religious faith. These items are not limited to consumer goods, and not all need to be traded in the market. To facilitate the subsequent discussion, assume that (5) takes the form of a nested CES, with an upper-tier of K broad categories of goods, services, actions and beliefs, and a lower tier of I_k varieties of each of these broad categories k, where $\sum_{k=1}^K I_k = I$:

$$u^j = \left(\sum_{k=1}^{K} \beta_k^j \left(\sum_{i=1}^{I_k} \alpha_{i_k}^j (x_{i_k}^j)^{\rho_v} \right)^{\frac{\rho_g}{\rho_v}} \right)^{\frac{1}{\rho_g}}. \tag{6}$$

In the above expression, $1/(1-\rho_g)$ is the elasticity of substitution between different categories (food, sports, spirituality, travel, family, etc.), and $1/(1-\rho_v)$ is the elasticity of substitution between different varieties of these categories (pizza, tacos, soccer, baseball, prayer, meditation, etc.). The parameters β_k^j capture how much individual j likes the different categories k and the parameters $\alpha_{i_k}^j$ capture how much she likes the different varieties of these categories. For convenience, we normalize these parameters by setting $\sum_k \left(\beta_k^j\right)^{\frac{1}{1-\rho_g}} = 1$ and $\sum_i (\alpha_{i_k}^j)^{\frac{1}{1-\rho_v}} = 1$.

By tracking the activity of individual j, Facebook receives signals of β_k^j and $\alpha_{i_k}^j$. One way of conceptualizing these signals is to view them as related to an individual's time use. While on Facebook (or one of the platforms Facebook has access to), assume that each individual is endowed with one unit of time that she allocates to different things (pizza, tacos, soccer, baseball, meditation, children's songs, …) in order to maximize:

$$\hat{u}^j = \left(\sum_{k=1}^{K} \hat{\beta}_k^j \left(\sum_{i=1}^{I_k} \hat{\alpha}_{i_k}^j (t_{i_k}^j)^{\rho_v} \right)^{\frac{\rho_g}{\rho_v}} \right)^{\frac{1}{\rho_g}} \tag{7}$$

where $t_{i_k}^j$ is the time individual j spends on i_k. The parameters $\hat{\beta}_k^j$ and $\hat{\alpha}_{i_k}^j$ are imperfectly related to the parameters β_k^j and $\alpha_{i_k}^j$ in (6). More specifically, assume that $\hat{\beta}_k^j = \beta_k^j + \varepsilon_k^j$ and $\hat{\alpha}_{i_k}^j = \alpha_{i_k}^j + \varepsilon_{i_k}^j$, where on average ε_k^j and $\varepsilon_{i_k}^j$ are zero. The underlying assumption is that users' online preferences (7) imperfectly reflect their offline preferences (6).[21] The time constraint of individual j can be written as:

$$\sum_k \sum_i (t_{i_k})^j = 1 \tag{8}$$

where the implicit price of spending one minute is the same for all i_k, and normalized to one. Maximizing (7) subject to (8) yields the time individual j spends on variety i_k:

$$t_{i_k}^j = \frac{(\hat{\alpha}^j)_{i_k}^{\frac{1}{1-\rho_v}}}{\sum_{i_k=1}^{I_k} (\hat{\alpha}^j)_{i_k}^{\frac{1}{1-\rho_v}}} \frac{(\hat{\beta}^j)_k^{\frac{1}{1-\rho_g}}}{\sum_{k=1}^{K} (\hat{\beta}^j)_k^{\frac{1}{1-\rho_g}}} = (\hat{\alpha}^j)_{i_k}^{\frac{1}{1-\rho_v}} (\hat{\beta}^j)_k^{\frac{1}{1-\rho_g}} \tag{9}$$

and the time individual j spends on category k:

$$t_k^j = \frac{(\hat{\beta}^j)_k^{\frac{1}{1-\rho_g}}}{\sum_{k=1}^{K} (\hat{\beta}^j)_k^{\frac{1}{1-\rho_g}}} = (\hat{\beta}^j)_k^{\frac{1}{1-\rho_g}} \tag{10}$$

[21] In this conceptual framework, we make a clear distinction between online activities that enter into (7) and offline activities that enter into (6). In a more complex model, this distinction would be less clear. On the one hand, time spent on Facebook would also enter into (6), and on the other hand, some offline activities would enter into (7) because Facebook has some ability to follow its users' whereabouts through GPS tracking.

where we normalize $\sum_k \left(\hat{\beta}^j_k\right)^{\frac{1}{1-\rho_g}} = 1$ and $\sum_i (\hat{\alpha}^j_{i_k})^{\frac{1}{1-\rho_v}} = 1$.

By observing an individual's time use patterns $t^j_{i_k}$ and t^j_k, Facebook uses (9) and (10) to estimate parameters $\hat{\alpha}^j_{i_k}$ and $\hat{\beta}^j_k$ in (7). It then uses this information to get binary estimates of parameters $\alpha^j_{i_k}$ and β^j_k in (6). Denote these binary estimates by $\mathbb{a}^j_{i_k}$ and $\mathbb{b}^j_k$. Facebook focuses on a binary, rather than a continuous, version of $\alpha^j_{i_k}$ and β_k because it aims to identify whether or not individual j likes a particular variety i_k or category k.[22] More specifically, if $\hat{\alpha}^j_{i_k}$ is above a certain threshold $\bar{\alpha}_{i_k}$, Facebook will conclude that j likes i_k:

$$\mathbb{a}^j_{i_k} = \begin{cases} 1 & : \quad \hat{\alpha}^j_{i_k} > \bar{\alpha}_{i_k} \\ & \\ 0 & : \quad \text{otherwise.} \end{cases}$$

That is, the binary variable $\mathbb{a}^j_{i_k}$ takes a value of 1 if $\hat{\alpha}^j_{i_k}$ is high enough, and it takes a value of 0 otherwise. The threshold $\bar{\alpha}_{i_k}$ depends both on the structure of ε_{i_k} and on how Facebook chooses to discretize its estimate. Similarly, Facebook concludes that an individual j likes category k if $\hat{\beta}^j_k$ is above a certain threshold $\bar{\beta}_k$:

$$\mathbb{b}^j_k = \begin{cases} 1 & : \quad \hat{\beta}^j_k > \bar{\beta}_k \\ & \\ 0 & : \quad \text{otherwise.} \end{cases}$$

The bottomline is that $\mathbb{a}^j_{i_k}$ and $\mathbb{b}^j_k$ are measures of user's j preferences for variety i_k and category k.

When observing many individuals of a given demographic group (say, women), Facebook can get an estimate of the number of female users for whom $\mathbb{a}_{i_k} = 1$ or $\mathbb{b}_k = 1$. For example, the number of female users with $\mathbb{a}_{i_k} = 1$ can be written as:

$$f^w_{i_k} = \sum_{j \in W} \mathbb{a}^j_{i_k}. \tag{11}$$

Likewise, the share of female users with $\mathbb{b}_k = 1$ can be written as:

$$f^w_k = \sum_{j \in W} \mathbb{b}^j_k. \tag{12}$$

The estimates f^w_{ik} and f^w_k can then be interpreted as the number of women who have revealed a preference for i_k and k.

Of course, Facebook data will only provide a good measure of the preferences of men and women across the globe to the extent that it is able to reliably identify a broad and comparable set of interests across countries and genders. We address these issues in further detail in the next section that discusses the data we use.

Online Appendix. Data sources and supplementary tables and figures

Details about the data sources and additional tables and figures can be found online at https://doi.org/10.1016/j.jebo.2025.107165.

Data availability

Data will be made available on request.

References

Abraham, L., 2023. The gender gap in performance reviews. J. Econ. Behav. Organ. 214, 459–492.
Aldén, L., Neuman, E., 2022. Culture and the gender gap in choice of major: An analysis using sibling comparisons. J. Econ. Behav. Organ. 201, 346–373.
Anghel, B., Rodríguez-Planas, N., Sanz-de Galdeano, A., 2020. Is the math gender gap associated with gender equality? Only in low-income countries. Econ. Educ. Rev. 79, 102064.
Archer, J., 2006. Testosterone and human behavior: An evaluation of the challenge hypothesis. Neurosci. Biobehav. Rev. 30 (3), 319–345.
Atari, M., Lai, M.H.C., Dehghani, M., 2020. Sex differences in moral judgements across 67 countries. Proc. R. Soc. Biol. Sci. 287, 20201201.
Baeza-Yates, R., Ribeiro-Neto, B., 2011. Modern information retrieval. In: The Concepts and Technology behind Search. Pearson Education Limited, Essex.
Balducci, M., 2023. Linking gender differences with gender equality: A systematic-narrative literature review of basic skills and personality. Front. Psychol. 14, 1105234.
Berenbaum, S., Beltz, A.M., 2021. Evidence and implications from a natural experiment of prenatal androgen effects on gendered behavior. Curr. Dir. Psychol. Sci. 30, 202–210.
Bertrand, M., 2011. New perspectives on gender. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, first ed. Vol. 4, Elsevier, pp. 1543–1590, chapter 17.
Boyd, R., Richerson, P.J., 2005. The Origin and Evolution of Culture. Oxford University Press, Oxford, U.K..
Breda, T., Jouini, E., Napp, C., Thebault, G., 2020. Gender stereotypes can explain the gender-equality paradox. Proc. Natl. Acad. Sci. 117 (49), 31063–31069.

[22] In advertising, a target audience requires a binary criterion that determines whether an individual belongs to the audience or not.

Buser, T., 2012. The impact of the menstrual cycle and hormonal contraceptives on competitiveness. J. Econ. Behav. Organ. 83 (1), 1–10.
Buss, D., 1989. Sex differences in human mate preferences: Evolutionary hypotheses tested in 37 cultures. Behav. Brain Sci. 12 (1), 1–14.
Cabañas, J.G., Cuevas, A., Arrate, A., Cuevas, R., 2020. Does Facebook use sensitive data for advertising purposes? Commun. ACM 64 (1), 62–69.
Caravaca, F., González-Cabañas, J., Cuevas, Á., Cuevas, R., 2024. Overprofiling analysis on major internet players. Proc. Priv. Enhanc. Technol. 4, 929–946.
Cárdenas, J.C., Dreber, A., Von Essen, E., Ranehill, E., 2012. Gender differences in competitiveness and risk taking: Comparing children in Colombia and Sweden. J. Econ. Behav. Organ. 83 (1), 11–23.
Carrer, L., De Masi, L., 2024. Municipal-Level Gender Norms: Measurement and Effects on Women in Politics. unpublished manuscript.
Charles, M., 2011. A world of difference: International trends in women's economic status. Annu. Rev. Sociol. 37 (1), 355–371.
Charles, M., Bradley, K., 2009. Indulging our gendered selves? Sex segregation by field of study in 44 countries. Am. J. Sociol. 114 (4), 924–976.
Cheryan, S., Ziegler, S.A., Montoya, A.K., Jiang, L., 2017. Why are some STEM fields more gender balanced than others? Psychol. Bull. 143 (1), 1–35.
Clots-Figueras, I., 2012. Are female leaders good for education? Evidence from India. Am. Econ. J. Appl. Econ. 4 (1), 212–244.
Costa, P.T., Terracciano, A., McCrae, R.R., 2001. Gender differences in personality traits across cultures: Robust and surprising findings. J. Pers. Soc. Psychol. 81 (2), 322–331.
Cotter, D., Hermsen, J.M., Vanneman, R., 2011. The end of the gender revolution? Gender role attitudes from 1977 to 2008. Am. J. Sociol. 117 (1), 259–289.
Croson, R., Gneezy, U., 2009. Gender differences in preferences. J. Econ. Lit. 47 (2), 448–474.
Cuberes, D., Teignier, M., 2014. Gender inequality and economic growth: A critical review. J. Int. Dev. 26 (2), 260–276.
Deerwester, S., Dumais, S.T., Furnas, G.W., Landauer, T.K., Harshman, R., 1990. Indexing by latent semantic analysis. J. Am. Soc. Inf. Sci. 41 (6), 391–407.
Dubois, A., Zagheni, E., Garimella, K., Weber, I., 2018. Studying migrant assimilation through Facebook interests. In: Staab, S., Koltsova, O., Ignatov, D. (Eds.), Social Informatics. SocInfo 2018, In: Lecture Notes in Computer Science, vol. 11186, Springer.
Eagly, A.H., Wood, W., Diekman, A.B., 2000. Social role theory of sex differences and similarities: A current appraisal. In: Eckes, T., Trautner, H.M. (Eds.), The Developmental Social Psychology of Gender. Erlbaum.
Eckart, C., Young, G., 1936. The approximation of one matrix by another of lower rank. Psychom. 1 (3), 211–218.
England, P., 2010. The gender revolution: Uneven and stalled. Gend. Soc. 24 (2), 149–166.
Englehardt, S., Narayanan, A., 2016. Online tracking: A 1-Million-site measurement and analysis. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 1388–1401.
Erikson, E.H., Erikson, J.M., 1998. The Life Cycle Completed (Extended Version). WW Norton & Company.
Falk, A., Hermle, J., 2018. Relationship of gender differences in preferences to economic development and gender equality. Sci. 362 (6412), eaas9899.
Fatehkia, M., O'Brien, D., Weber, I., 2019. Correlated impulses: Using Facebook interests to improve predictions of crime rates in urban areas. PLoS One 14 (2), e0211350.
Fernández, R., 2014. Women's rights and development. J. Econ. Growth 19 (1), 37–80.
Fors Connolly, F., Goossen, M., Hjerm, M., 2020. Does gender equality cause gender differences in values? Reassessing the gender-equality-personality paradox. Sex Roles 83, 101–113.
Funk, P., Gathmann, C., 2015. Gender gaps in policy making: Evidence from direct democracy in Switzerland. Econ. Policy 30 (81), 141–181.
Gächter, S., Schulz, J.F., 2016. Intrinsic honesty and the prevalence of rule violations across societies. Nat. 531 (7595), 496–499.
García, D., Mitike Kassa, Y., Cuevas, A., Cebrian, M., Moro, E., Rahwan, I., R. Cuevas, R., 2018. Analyzing gender inequality through large-scale Facebook advertising data. Proc. Natl. Acad. Sci. 115 (27), 6958–6963.
Gneezy, U., Leonard, K.L., List, J.A., 2008. Gender differences in competition: Evidence from a matrilineal and patriarchal society. Econom. 77 (3), 909–931.
Guiso, L., Monte, F., Sapienza, S., Zingales, L., 2008. Culture, gender, and math. Sci. 320 (5880), 1164–1165.
Herlitz, A., Hönig, I., Hedebrant, K., Asperholm, M., 2025. A systematic review and new analyses of the gender-equality paradox. Perspect. Psychol. Sci. 20 (3), 503–539.
Holmlund, H., Rainer, H., Reich, P., 2023. All geared towards success? Cultural origins of gender gaps in student achievement. J. Econ. Behav. Organ. 209, 222–242.
Kaiser, T., 2019. Nature and evoked culture: Sex differences in personality are uniquely correlated with ecological stress. Pers. Individ. Differ. 148, 67–72.
Kosakowska-Berezecka, N., Besta, T., Jurek, P., Olech, M., Sobiecki, J., Bosson, J., Zadkowska, M., 2024. Towards gender Harmony dataset: Gender beliefs and gender stereotypes in 62 countries. Sci. Data 11 (1), 392.
Kosinski, M., Stillwell, D., Graepel, T., 2013. Private traits and attributes are predictable from digital records of human behavior. Proc. Natl. Acad. Sci. 110 (15), 5802–5805.
Kuhn, A., Wolter, S.C., 2022. Things versus people: Gender differences in vocational interests and in occupational preferences. J. Econ. Behav. Organ. 203, 210–234.
Levanon, A., Grusky, D.B., D. B., 2016. The persistence of extreme gender segregation in the twenty-first century. Am. J. Sociol. 122 (2), 573–619.
Liberini, F., Redoano, M., Russo, A., Cuevas, A., Cuevas, R., 2020. Politics in the Facebook Era - evidence from the 2016 US presidential elections. In: CESifo Working Paper Series. Vol. 8235.
Lippa, R.A., Collaer, M.L., Peters, M., 2010. Sex differences in mental rotation and line angle judgments are positively associated with gender equality and economic development across 53 nations. Arch. Sex. Behav. 39 (4), 990–997.
Lippmann, Q., 2021. Gender and lawmaking in times of quotas. In: Working Paper.
Mac Giolla, E., Kajonius, P.J., 2019. Sex differences in personality are larger in gender equal countries: Replicating and extending a surprising finding. Int. J. Psychol. 54 (6), 705–711, 2019.
Mejova, Y., Gandhi, H.R., Rafaliya, T.J., Sitapara, M.R., Kashyap, R., Weber, I., 2018. Measuring subnational digital gender inequality in India through gender gaps in Facebook use. In: COMPASS '18 Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies. Vol. 43, pp. 1–5.
Molina, T., Usui, E., 2023. Female labor market conditions and gender gaps in aspirations. J. Econ. Behav. Organ. 211, 165–187.
Napp, C., Breda, T., 2022. The stereotype that girls lack talent: A worldwide investigation. Sci. Adv. 8 (10), eabm3689.
Neyse, L., Johannesson, M., Dreber, A., 2021. 2D:4D does not predict economic preferences: Evidence from a large, representative sample. J. Econ. Behav. Organ. 185 (C), 390–401.
Niederle, M., Vesterlund, L., 2011. Gender and competition. Annu. Rev. Econ. 3 (1), 601–630.
Obradovich, N., Özak, Ö., Martín, I., Ortuño Ortín, I., Awad, E., Cebrián, M., Cuevas, R., Desmet, K., Rahwan, I., Cuevas, A., 2022. Expanding the measurement of culture with a sample of two billion humans. J. R. Soc. Interface 19, 20220085.
Pew Research Center, 2019. Facebook algorithms and personal data.
Quisimbing, A.R., Maluccio, J.A., 2020. Intrahousehold allocation and gender relations: New empirical evidence from four developing countries. In: FCND Discussion Paper 84. International Food Policy Research Institute (IFPRI).
Rama, D., Mejova, Y., Tizzoni, M., Kalimeri, K., Weber, I., 2020. Facebook ads as a demographic tool to measure the urban-rural divide. In: Proceedings of the Web Conference 2020. WWW '20, Taipei, Taiwan, April 20-24.
Richardson, S., 2020. Gender stereotypes, gendered self-expression, and gender segregation in fields of study: A Q & A with professor maria charles. GenderSci Blog. February 14. GenderSciWebLink.
Samek, A., 2019. Gender differences in job entry decisions: A university-wide field experiment. Manag. Sci. 65 (7), 3272–3281.

Sapienza, P., Zingales, L., Maestripieri, D., 2009. Gender differences in financial risk aversion and career choices are affected by testosterone. Proc. Natl. Acad. Sci. 106 (36), 15268–15273.
Schmitt, D.P., 2015. The evolution of culturally-variable sex differences: Men and women are not always different, but when they are... It appears not to result from patriarchy or sex role socialization. In: Shackelford, T.K., Hansen, R.D. (Eds.), The Evolution of Sexuality. Springer International Publishing.
Schmitt, D.P., Long, A.E., McPhearson, A., O'Brien, K., Remmert, B., Shah, S.H., 2017. Personality and gender differences in global perspective. Int. J. Psychol. 52, 45–56.
Shores, T.S., 2007. Applied Linear Algebra and Matrix Analysis. Springer, New Springer, New York.
Stoet, G., Geary, D.C., 2018. The gender-equality paradox in science, technology, engineering, and mathematics education. Psychol. Sci. 29 (4), 581–593.
Su, R., Rounds, J., Armstrong, P.I., 2009. Men and things, women and people: A meta-analysis of sex differences in interests. Psychol. Bull. 135 (6), 859–884.
Vieira, C.C., Vasconcelos, M., 2021. Using Facebook ads data to assess gender balance in STEM: Evidence from Brazil. In: Proceedings of the Web Conference. 2021.
Vishkin, A., 2022. Queen's gambit declined: The gender-equality paradox in chess participation across 160 countries. Psychol. Sci. 33 (2), 276–284.
Wood, W., Eagly, A.H., 2012. Biosocial construction of sex differences and similarities in behavior. In: Olson, J.M., Zanna, M.P. (Eds.), Advances in Experimental Social Psychology. Vol. 46, Academic Press, pp. 55–123.
Wozniak, D., Harbaugh, W.T., Mayr, U., 2014. The menstrual cycle and performance feedback alter gender differences in competitive choices. J. Labor Econ. 32 (1), 161–198.
Youyou, W., Kosinski, M., Stillwell, D., 2015. Computer-based personality judgments are more accurate than those made by humans. Proc. Natl. Acad. Sci. 112 (4), 1036–1040.