

## The Impact of the COVID-19 Pandemic on Women's Contribution to Public Code

Annalí Casanueva · Davide Rossi ·  
Stefano Zacchirolì · Théo Zimmermann

**Abstract** Despite its promise of openness and inclusiveness, the development of free and open source software (FOSS) remains significantly unbalanced in terms of gender representation among contributors. To assist open source project maintainers and communities in addressing this imbalance, it is crucial to understand the *causes* of this inequality.

In this study, we aim to establish how the COVID-19 pandemic has influenced the ability of women to contribute to public code. To do so, we use the Software Heritage archive, which holds the largest dataset of commits to public code, and the difference in differences (DID) methodology from econometrics that enables the derivation of causality from historical data.

Our findings show that the COVID-19 pandemic has disproportionately impacted women's ability to contribute to the development of public code, relatively to men. Further, our observations of specific contributor subgroups indicate that COVID-19 particularly affected women hobbyists, identified using contribution patterns and email address domains.

---

A. Casanueva  
Ifo institute, Big Data Junior Research Group, Munich, Germany.  
E-mail: casanuevaartis@ifo.de

D. Rossi  
University of Bologna, Bologna, Italy.  
E-mail: daviderossi@unibo.it

S. Zacchirolì  
LTCI, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France.  
E-mail: stefano.zacchirolì@telecom-paris.fr

T. Zimmermann  
LTCI, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France.  
E-mail: theo.zimmermann@telecom-paris.fr

## 1 Introduction

Women have historically been, and remain to this day, underrepresented in free/open source software (FOSS) development [52]. The phenomenon is not uncommon in broader contexts like STEM (Science, Technology, Engineering, and Mathematics) disciplines [9, 12, 29, 43, 53] and computing [24, 34], but it is particularly severe in FOSS, where women contributions have repeatedly been observed to be as low as 5–10% [52], in terms of both numbers of commits and active participants. Trends over the past decade have been moving towards increased women participation in FOSS [54]—and, more generally, in *public code*, to refer to all code developed on public collaborative development platforms, no matter what the software license is—but the bottom line remains tilted towards very low women participation.

Both the negative effects of such a gender gap and needed actions to counter it have been explored in previous work; see Trinkenreich et al. [51] for a high-level overview of where we stand at the time of writing. Generally speaking, actions to counter the gender gap and encourage women participation in FOSS can happen at different levels: locally within specific FOSS communities (e.g., adopting a code of conduct to prevent harassment) and more globally (e.g., programs like Outreachy or Girls Can Code as well as state-level policies). No matter the policy level, actions need to be based on a solid understanding of the *causes* of the low participation of women in FOSS.

With this study we aim to contribute to such understanding by exploring the impact of the COVID-19 pandemic—which started in early 2020 and led to lockdowns and other stringency measures around the world—on women contributions to public code.

Although major shocks like the COVID-19 pandemic are beyond the control of individual FOSS projects, they can still have a significant impact on women’s participation in these projects. It is therefore essential that FOSS projects that take measures to increase the inclusion of women are aware of the impact of such external shocks, and *can quantify their effect*, as part of a more encompassing monitoring of the effectiveness of their initiatives. Furthermore, by leveraging their comprehension of external factors contributing to inequality, FOSS projects can foster a culture that acknowledges and accounts for these dynamics, e.g., when implementing meritocratic principles to determine who to promote in the community. At the policy level, understanding the impact of the COVID-19 pandemic or similar future shocks will allow directing funds and actions towards the most effective interventions.

In fields other than computing, it has been observed [3, 13, 15, 16, 21, 25, 37, 42, 49, 55] (see Section 2.1 for a detailed breakdown of these related works) how COVID-19 correlated with and/or heightened gender disparities, due to both childcare responsibilities (that tend to fall more on mothers’ shoulders) and workplace dynamics. In computing, the impact of COVID-19 on software development has been explored [5, 14, 23, 31, 35, 36, 41, 46] (see Section 2.2 for a detailed breakdown), focusing for the most part on the effects of the pandemics on developer productivity, work patterns, and well-being. For public

code contributions, a recent study [44] observed a first-time decrease in women contributions in 2020 and interpreted it as possibly explained by the COVID-19 pandemic, without attempting to properly establish a causal relationship between the two. It is hence still unknown whether the arrival of the COVID-19 pandemic *directly caused* an impairment in the ability of women to contribute to public code, relatively to men.

### 1.1 On the importance of establishing causality

In this work, we aim at going beyond observing correlations. Documenting a correlation differs from identifying a causal relationship and taking one for the other could lead to a false interpretation of the situation and ultimately to creating ineffective (if not harmful) policies.

Let us consider a simplified example that could apply to our case of study. Imagine the world is composed of two countries: country A and country B. In country A, the state has not enough budget to provide daycare facilities for children or dependent persons. As a consequence, care responsibilities are handled by the family and, in particular, by women. These responsibilities take time and impede women’s contributions to public code. Country B, on the other hand, has a good state budget and is able to provide public services that take the burden of care responsibilities. In country B, women have more time to contribute to public code. When the COVID-19 pandemic hits, country B is able to provide people with subsidies that allow them to stop working and make it easier to respect the social distancing measures. On the contrary, country A has no budget to provide subsidies, people have a higher need to work and as a result are less able to avoid close interactions. This leads to a higher incidence of COVID-19 in country A.

Looking at the data of these two countries, we will find a correlation between higher incidence of COVID-19 and lower contribution of women to public code. In this case, a naive observer could think that the COVID-19 pandemic was the cause of the lower contribution of women to public code, when in fact there is a third factor (a hidden variable or confounder) that is the real cause of both the higher incidence of COVID-19 and the lower contribution of women to public code: state budget.

Now consider a private foundation dedicated to supporting open source, and which has as one of its missions to foster gender diversity in open source projects. When the next pandemic hits, this foundation, basing its decisions on the naive conclusion from observed correlations, wants to avoid a negative impact on gender diversity in public code, and they may be inclined to redirect some of their funding to distributing masks to the public. In this fictional example, this strategy will turn out to be completely ineffective for its objective, since the real cause of the lower contribution of women was state budget, and not the pandemic. Had they known the real cause of the lower contribution of women, they could have used their resources more effectively, e.g., by supporting measures that relieve women from care responsibilities.

Beyond identifying actual *causes* affecting women’s ability to contribute to public code, *measuring the impact* of each cause is also key. When resources for an intervention are limited, knowing which factors have a greater impact, allows to allocate resources more effectively. Alternatively, when several factors play a role simultaneously, knowing the impact of each one allows to determine if a specific intervention is effective or not and permits to truly address the problem at hand.

Imagine for example that a well-intentioned group of software developers implemented a code of conduct in the community with the goal that women feel more welcome and are able to contribute more. They implemented this code of conduct during the pandemic. To evaluate if their intervention had an effect, they measured the relative level of activity of women before and some weeks after the intervention. They found out, to their deception, that after the implementation of their new code of conduct, women in fact contributed relatively *less* than before. Being able to factor out the direct effect of COVID-19 would have made them realize that in fact their policy was *effective* as the decrease in the relative level of contribution was in fact lower than what would have happened without the intervention. The culprit to the observed decrease in relative contribution is COVID-19 and the code of conduct was, in fact, able to counter-balance its effect a little.

Building a body of empirical literature that is able to *identify the causes* that prevent women from contributing to public code in larger numbers and *quantitatively measure* them is key, particularly in a field that aims to identify interventions that could help to effectively increase diversity.

## 1.2 Contributions

*Research questions.* This work explores the general theme of: what were the effects of the COVID-19 pandemic on gender diversity in contributions to public code? More precisely, we will answer the following research questions:

- **RQ1** *How did the COVID-19 pandemic impact the ability of women (relatively to men) to contribute to public code?*
- **RQ2** *Which groups of women contributing to public code were impacted the most (relatively to men in the same groups), in their ability to contribute, by the COVID-19 pandemic?*

*Methodology.* In order to answer these research questions, we start from a large-scale dataset obtained from the Software Heritage archive [38] and comprised of 3.8 B (billion) commits contributed by 71 M (million) authors via major software development forges (e.g., GitHub, GitLab.com) and source code distribution platforms (e.g., package manager repositories), over a period of several decades. We then, on the one hand, detect the gender of contributors using name-based heuristics and, on the other hand, geolocate the same population to individual countries around the world by applying heuristics

based on the domain of contributor emails as recorded in commit data. (See Section 3 for details.)

Data obtained this way allows observing commits to public code before and during the pandemic and to aggregate them by country, period, and contributor groups. It also allows studying the track record of contributors via features like: when they started contributing, in which years they were active, and other time-based work patterns.

We complement large-scale public code data with world population and COVID-19 data, and most notably: worldwide daily COVID-19 deaths reported to WHO (World Health Organization), global population data from the UN (United Nations), and the Oxford COVID-19 dataset about the level of stringency of governments responses to COVID-19.

Then, we answer the stated research questions applying *difference in differences* (DID) [4], a technique from econometrics that allows us to unveil causality, rather than mere correlation, between observed changes in selected variables, such as being a woman or having contributed to public code in a given period and living in a country where people were more exposed to COVID-19 restrictions. (See Section 4 for details.)

*Key findings.* We answer RQ1 by showing that higher exposure to COVID-19 caused a statistically significant reduction in contributing activity (as measured in number of commits per week) by women, relatively to men. The result is robust and holds for two different measures of exposure to COVID-19: COVID-19 deaths *per capita* in a given country and strictness of anti-COVID-19 measures. For the latter COVID-19 measure, our result also holds when changing the measure of contributing activity from number of commits to number of active days (with at least one commit) per week. In particular, one additional COVID-19 related death per 100 000 inhabitants widens the gap of the number of contributions between men and women by 2.1%. Alternatively, passing from a situation with no restrictions to a situation with social distance measures of 75% of the maximum of our measure, increases the gap by around 28%.

Regarding RQ2, we explore four different ways of categorizing contributors. Two of these categorizations aim to distinguish between hobbyists and paid contributors. Both of them lead us to the same conclusion: the relative impact of COVID-19 on women is stronger for hobbyists, as identified by contributing mostly outside of working hours and by using non-professional email addresses.

### 1.3 Paper structure

Section 2 reviews related work on COVID-19, gender issues, and software development. Section 3 presents the datasets we used and how we treated and augmented them, including gender detection and geolocation. The use of difference in differences is discussed in Section 4. Section 5 presents obtained results, answering the stated research questions. The interpretation of our

findings and threats to their validity are discussed in Section 6. Section 7 concludes the paper and suggests leads for future work.

#### 1.4 Data availability

A replication package for the work presented in this paper is available [11]. See Section 6.2.3 for more information about what it includes.

## 2 Related work

The effects of the COVID-19 pandemic on various organizations have been extensively investigated and analyzed, delving into the multifaceted ways in which these events have influenced the dynamics of involved entities [19, 30]. The present work sits at the crossroad between: (1) gender issues, (2) the COVID-19 pandemic, and (3) software engineering in society. While these three topics *together* have only been studied tangentially (see [32] for an exception), various intersections among them have been analyzed. In this section, we review numerous related works in this broad space, organizing them based on the possible overlaps of the three topics.

### 2.1 COVID-19 and gender issues

A first body of related work looked at the potential of the COVID-19 crisis to amplify gender disparities in both workplace dynamics and household management tasks (houseworks, children and elderly care). Most of these studies operate on data elicited via online surveys and adopt regression analysis to assess the correlation between the pandemic and gender disparities. Apart from some minor methodological differences, what differentiates these studies is the geographical origin of the respondents. Most studies focus on a specific nation/geographic area, with most findings suggesting that emerging phenomena are common across America, Europe, and Japan.

*Takeaways.* Existing studies on the impact of COVID-19 on the gender gap, reviewed in details below, establish correlation rather than causation (with a few exceptions, like [25]) between the two. This marks a major difference with respect to this paper, where we employ difference in differences (DID) to establish that changes in public code contributions by women are indeed *due to* COVID-19.

*Study details.* Adams-Prassal et al. [3] use real-time survey data, gathered in the US, UK, and Germany in April 2020, to examine changes in income, employment, mental health, and care work among different groups of people during the pandemic. The paper highlights that while everyone has been

affected by COVID-19, there are significant differences across socioeconomic groups. Authors also find that women have been disproportionately affected by the pandemic due to their higher care burden, both before and during the crisis. For example, they found that in the US and the UK, women were 6.5% and 4.8% more likely to lose their jobs.

Collins et al. [13] use panel data from the US Current Population Survey to examine changes in mothers' and fathers' work hours from February through April 2020. Using fixed effects models they found that mothers with young children have reduced their work hours 4–5 times more than fathers.

Del Boca et al. [16] employ survey data gathered in April 2020 from a representative sample of Italian women. It seeks to examine the effects of altered working arrangements due to COVID-19 on housework, childcare, and homeschooling within couples where both partners work. Authors discovered that, except for women working at their usual workplace, all surveyed women spent more time on housework. Men's housework time depends on their partners' working arrangements, increasing when partners work away. Childcare time is symmetrically affected, with both women and men spending less time if working away. Working women with young children find work-life balance challenging during COVID-19, especially if their partners continue to work outside the home.

Sevilla et al. [49] present a regression analysis of 2782 responses to the Ipsos MORI omnibus survey of May 2020, provided by respondents living in the UK, aged 18–60, part of families with children under the age of 12. The authors found that the increase in childcare hours for women is less influenced by their employment status compared to men. This often results in many women managing both work and a significantly increased childcare burden, potentially leading to adverse effects on their mental health and future career prospects. Notably, households where men were not employed show greater strides toward achieving a more equal distribution of childcare responsibilities.

Zamarro et al. [55] analyzed approximately 7000 respondents to the Understanding Coronavirus in America Tracking Survey in March 2020 to understand how fathers and mothers were coping with this crisis in terms of childcare provision, employment, working arrangements, and psychological distress levels. They found that women shouldered a greater childcare burden than men during the COVID-19 crisis, even while working. The current working situations of mothers seem to have minimal impact on their childcare responsibilities. However, this division of childcare is linked to reduced working hours and an increased likelihood of working mothers transitioning out of employment.

Dang et al. [15] analyzed 6089 responses to surveys collected across 6 countries including China, South Korea, Japan, Italy, the United Kingdom and the four largest states in the United States, allowing an analysis that is not specific to a single country's dynamics. The survey contained questions on fundamental demographic factors of participants, their employment and living arrangements, health status and diseases, self-reported economic and non-economic consequences of the pandemic, behaviors, beliefs about the pandemic, and evaluations of government responses (parts of a dataset described in [6]). The

regression analysis conducted by the authors indicate that women face a 24 percent higher likelihood of experiencing permanent job loss due to the outbreak compared to men. Additionally, women anticipate a 50 percent greater reduction in their labor income compared to men.

Farre et al. [21] analyzed 5001 responses to a survey distributed in Spain to a population aged 24-50 in May 2000. The questions focused on changes in employment during the lockdown and changes in the distribution of childcare and housework. The authors' analysis reveals that nearly 20% of workers, particularly those with lower education levels, experienced furloughs. Temporary job losses slightly affected women more than men, and women were more likely to work from home during confinement. The study also found a significant increase in childcare and housework responsibilities for parents, with women continuing to bear the majority of the burden. On average, there is a gender gap of about 17% in parents' shares of childcare and housework during the lockdown.

Reichelt et al. [42] collected 5008 responses from a survey conducted in May–June 2020 in the USA, Germany and Singapore. The panel included individuals working either full- or part-time in January 2020, providing retrospective responses about their employment status and couples where both the respondent and their cohabiting partner worked full- or part-time in January, offering retrospective information about themselves and their partners. Their regression analyses indicate that women experienced more frequent transitions to unemployment, reductions in working hours, and shifts to working from home compared to men, although the extent varied across the three countries studied. Additionally, among couples initially employed at the start of the pandemic, men tended to express more egalitarian gender-role attitudes when they became unemployed while their partners remained employed. In contrast, women expressed more traditional attitudes when they became unemployed and their partners remained employed.

Petts et al. [37] used data from U.S. parents (623 mothers and 891 fathers) to empirically investigate whether the challenges posed by the loss of childcare and new homeschooling demands are linked to employment outcomes early in the pandemic. They also explore whether the pre-pandemic division of child care is associated with parents' employment. The authors find, by the means of regression studies, that for parents with young children, the loss of full-time childcare increased the risk of unemployment for mothers but not for fathers. However, greater father involvement in childcare helped mitigate negative employment outcomes for mothers of young children. In the case of parents with school-age children, participation in homeschooling was associated with adverse employment outcomes for mothers but not for fathers.

Fukai et al. [25] used data from the Labor Force Survey, published by Japan's Ministry of Internal Affairs and Communications, encompassing about 40 000 monthly responses dating back to well before the onset of the pandemic. This allowed the study to adopt a regression discontinuity design (RDD) approach and estimate that the pandemic decreased (causation) the mothers'

employment rate by about 4%, mostly because of increased childcare responsibilities.

Other relevant works [8,39] do not offer analyses of new data; instead, they discuss the results presented in previous works, aiming to uncover underlying societal issues and propose policies addressing gender inequality.

## 2.2 COVID-19 and software development

Professional software development teams, as well as online open source communities, also attracted the attention of researchers. A second body of relevant literature follows the tradition of empirical software engineering, utilizing both elicited and mined data, offering analyses that encompass both quantitative and qualitative aspects.

*Takeaways.* Overall, empirical studies on COVID-19 on software development, open source or otherwise, have found various sorts of correlations between the two, ranging from productivity changes to work patterns (work hours and weekdays) and from what activities are performed (e.g., development vs documentation writing) to developers’ well-being. We complement this body of work by showing that COVID-19 has caused, rather than merely being associated with, a decrease in the ability of women to contribute to public code, providing a detailed analysis of which groups of women from the pre-pandemic population were affected.

*Study details.* Ralph et al. [41] examined 2225 questionnaires to assess the effects of the pandemic on software developers. The study adopts non-parametric hypothesis tests and structural equation modeling to establish a negative correlation between developers working from home because of social restrictions and their well-being and productivity.

Klotzman et al. [31] used data mined from GitHub and Stack Overflow to determine if COVID-19 impacted open source development. The authors could not find conclusive indications about GitHub, but they found out that there has been a significant increase in the number of new users and questions posted on Stack Overflow during the first months of the pandemic (March/April 2020).

Ford et al. [23] analyzed a couple of online surveys, one with 1369 responses and another with 2265 responses, distributed to two sets of workers (the first in King County, WA, USA; the second in all USA states) among software engineers at Microsoft. The surveys were conducted from March to May 2020, when the company had support in place for its developers and remote work. The authors discuss aspects such as productivity, benefits, challenges, and recommendations that emerged from the surveys. They found a dichotomy in developer experiences influenced by various factors: what is perceived as a benefit for some is considered a challenge by others. For instance, the necessity of staying at home is seen as a benefit for being close to family

members but as a challenge for those struggling with working space and dealing with interruptions. The analysis of the responses not only addresses the different perceptions of the implications of the pandemic, but also provides deeper insights into the narratives stemming from these diverse experiences.

McDermott et al. [35] extracted data from the GitHub timeline to assess changes in work patterns potentially associated with the advent of the pandemic. Their work shows that developers were more active on weekends and outside of regular work hours than they were in previous years. It also detects that COVID-19 was associated with a jump in overall activity: during the early lockdown period, GitHub users were not merely reallocating a fixed budget of work hours; instead, they were actively increasing their overall work hours. Self-reported men and women users exhibited comparable patterns of increased working hours shortly after March 2020. However, men responded more promptly and consistently to the challenges posed by the pandemic. In contrast, women were slower in reallocating work, and their response was more inconsistent.

Russo et al. [46] describe a two-wave longitudinal study involving almost 200 selected software professionals on a global scale. Aim of the study was to infer pandemic-induced changes in everyday tasks correlated with perceived well-being, productivity, and other pertinent psychological and social factors. The findings presented by the authors indicate that the time allocated to specific activities from home was comparable to when working in the office. However, the duration developers devoted to each activity showed no correlation with their well-being, perceived productivity, and other variables.

Da Mota Silveira Neto et al. [14] combined a mining approach (conducted on 100 GitHub Java projects) with the analysis of 279 survey responses. Various code and project metrics were used to assess productivity, while the surveys aimed to evaluate projects and developers' well-being. In line with other previously discussed studies, the authors found that the impact of COVID-19 is not a strict dichotomy (reducing vs. increasing productivity); rather there exists a spectrum, where significant proportions of respondents hold differing opinions on the matter.

Pejić et al. [36] analyzed the GitHub timeline, comparing the event patterns in 2017-2019 and comparing them with the period 2020-2022. They observed that, overall, events associated with individual development have either maintained or experienced an increase in their trends. In contrast, events linked to community activity (such as forking) or contributions to documentation have exhibited a slight decrease.

Bao et al. [5] analyzed developers' daily activities (commits, code reviews, build, etc.) from Baidu Inc., one of China's largest IT companies. The authors aimed to evaluate the impact of the pandemic-mandated work-from-home arrangement on productivity. Similar to other studies, they identify both positive and negative impacts

### 2.3 Software development and gender issues

Leaving COVID-19 aside, a third and final body of literature have analyzed in recent years, gender-related issues within software development teams. A comprehensive reference that analyzes the most relevant works on this topic is Trinkenreich et al. [52], which provides a framework to position the various contributions and to which we refer the interested reader.

Usually studies operate on elicited data (mostly online surveys), mined data, or both. Those operating on mined data extract various types of items, using them as traces of productivity and/or collaboration: commits, issues, pull requests, code reviews, but also social software-like elements, such as GitHub “stars”.

While most of the presented works mainly aim at developers operating on FOSS/public code projects (all, among those adopting software mining techniques), there are notable exceptions. For example, [50] is a recent study based on the analysis of a survey (94 respondents) conducted among Ericsson’s female workforce, allowing the authors to identify and assess challenges (maternal wall, glass ceiling, impostor syndrome, etc.) and propose associated mitigations (support work-life balance, empower women, support women’s career growth, etc.).

*Takeaways.* The overall picture painted by these works is disheartening: under-representations and biases against women feed each other in an endless loop. The level of awareness, however, is rising; a few strategies to ease (and retain) women’s participation are becoming more common (like the very elementary adoption of explicit codes of conduct), and recent positive trends have been observed [40, 44], until the arrival of the pandemics.


The present work aims at completing the picture of the current state of gender issues in software development, by looking at the intersection of public code, COVID-19, and the causal relation among the two.

## 3 Dataset

To answer the stated research questions, we need data about public code production over time, as well as geolocated data about the COVID-19 pandemic. In this section, we describe the datasets we obtained to address these needs and how we preprocessed them before further analysis. The rest of our methodology pipeline, from the dataset on, is detailed in Section 4.

### 3.1 Public code data

*Initial dataset.* As our starting point for public code data we obtained from the Software Heritage initiative [38] a dump of all the commits archived by the project up to March 2024. Software Heritage [2] is the largest archive of



**Fig. 1** Public code data used in this study. Starting from commits archived by Software Heritage, we split them into author and commit data, classifying each class along two axes: gender and country of origin.

source code and its development history, as captured by modern version control systems (VCS), like Git. It archives both code and commit data from hundreds of million projects developed on, or distributed from, major development forges (GitHub, GitLab, Bitbucket, etc.), free/open source software (FOSS) distributions (Debian, GNU, Nix, etc.), and package manager repositories (PyPI, NPM, Maven, etc.). The specific data dump version we obtained contains 3 808 132 978 commits (unique by commit SHA1, computed in the same way Git does), authored by 70 778 755 authors (unique by full name strings, as recorded in VCSs). The dataset does not contain the actual files (source code, data, etc.) associated to each commit, but rather all *metadata*, such as authors, timestamps, commit messages, etc.

Obtained commits came as two relational tables, one for commits and one for authors. Each row in the commit table contains the following columns: commit SHA1 identifier, author and committer timestamps, author and committer identifiers. Each of the last two columns references the author table via a foreign key. The distinction between authors and committers comes from Git, which allows a code integrator to commit a change authored by someone else. For this study, we used authors and ignored committers, because the difference between the two is immaterial for our research questions and negligible in amount. Each row in the author table contains two fields: full name and email, parsed (by Software Heritage) from the raw author identification strings stored natively by VCSs in the form "**Author Name** <foo@example.com>".

From this point on, the preprocessing steps applied to public code data are those depicted in Figure 1 and described in the following. Some of them—Author cleanup and Gender detection—have been implemented, with specific deviations described below, by replaying the replication package of Rossi and Zacchiroli [44] on the obtained version of Software Heritage data. Crucially, this version spans 2.5 additional years of public code data (with respect to [44]), encompassing the arrival of the COVID-19 pandemic and accompanying stringency measures. However, we also learned that there are significant delays in the Software Heritage archiving of some public repositories, which could bias

the resulting data analysis,<sup>1</sup> therefore we only consider commits up to May 2022 in our analysis, for which we have exhaustivity guarantees despite the archiving delay. Other parts of the data processing pipeline were implemented from scratch, as detailed below.

*Data cleanup.* The *author cleanup* step removes from the initial dataset authors whose names are either unusable for practical reasons, or implausible as names of real people. Specifically, this step drops: names that cannot be decoded as UTF-8 strings, email addresses used instead of (or in addition to) names, names consisting of only blank characters, names containing more than 10% non-letter Unicode characters, and names longer than 100 characters. About one quarter of the initial author strings were removed by this filtering step, with 51 664 829 authors remaining for further analysis.

As it is customary for commit analysis practices, we further removed bots from the dataset (*bots removal* in Figure 1) and merged together authors appearing with different identities (*identity merging*), in this order. For bots removal, we used the approach and dataset of Dey et al. [17], removing from the dataset all authors whose emails match a bot in their dataset. After this step, 51 617 013 authors remained.

For identity merging, we implemented the straightforward approach of merging together all authors sharing the same email address. Also, when the same email address is used by ten or more authors we dropped them all. This was usually due to the adoption of generic, non realistic, email addresses like `you@example.com` (38K occurrences), `devnull@localhost` (12K occurrences), and so forth. After this step, 40 913 389 authors remained, referred to as *Unique human authors* in Figure 1.

*Gender detection.* To answer the stated research questions, we need to associate a gender and a country to the authors of commits in our dataset. Some preliminary considerations are in order before attacking the first task, which we refer to as “gender detection” in the following. At this scale (tens of million authors), manual approaches based on interviews, where individuals state what gender they identify with, are not feasible. Hence, following in the steps of previous work at similar scales [44,54], we rely instead on automated mechanisms (described later) that detect the gender of authors based on their public names, as recorded in code commits. Those tools are generally restricted to binary gender assignment. With that, we do not intend to arbitrarily define people within a binary gender confinement regardless of their preferences and sensitivity. In fact, none of the gender-related decisions made by the automated techniques used in this paper make sense when applied to *individuals* present in the analyzed corpus. The meaning of the exercise is statistical in nature and aims only to address the stated research questions. The used approach makes sense only in aggregate form and carries with it the unavoidable

---

<sup>1</sup> Active repositories are archived every two weeks on average. However, new repositories are archived in a different pipeline, which prioritizes recently updated repositories, leading to an archiving delay that can grow up to two years for new, but inactive, repositories.

limitations that name-based gender detection entail. We elaborate further on those limitations in Section 6.2.

To detect the gender of an author from our public code dataset, we apply `GENDER-GUESSER` [22] to its full name, as recorded in a commit. `GENDER-GUESSER` is an open source Python library for gender detection, based on first name frequencies around the world, which is often used in related work. Our choice of `GENDER-GUESSER` over alternatives is based on the fact that it is open source, which helps with study repeatability, and an actual self-hostable tool, rather than a remote service API which would be unusable (due to costs or execution time) at this scale. Furthermore, Santamaria et al. [48] conducted a comparative benchmark of `GENDER-GUESSER` and its main competitors, including payware commercial API services, showing that even if `GENDER-GUESSER` is not the best tool in absolute terms, it works comparatively well with geographically diverse datasets, as ours is by construction (it contains commits coming from all over the world).

As in previous works based on similar datasets and `GENDER-GUESSER` [44, 54], we address the issue that `GENDER-GUESSER` works primarily on first (given) names by applying a *majority criterion* for the final gender determination. We first tokenize full name strings by splitting at each blank, hyphen, or case change (to also address CamelCase notation, which is used by several authors in the dataset), and then use `GENDER-GUESSER` to determine the gender of each token. Among the tokens that were associated with a gender, if and only if a *strict majority* of name tokens for a given author full name is detected as belonging to one gender, we associate the majority gender to the author; otherwise their gender will remain unknown.

Differently from Rossi and Zacchiroli [44], due to identity merging we reach this data processing step with multiple author names, each composed by multiple tokens, associated to each unique author. To associate a gender to each *unique* author, we first applied the majority criterion token by token to each identity, and then applied it again on the genders obtained for each identity to determine the final gender of each unique author.

In the end, we identified a total of 13 354 546 men contributors and 2 398 896 women contributors, counted as unique authors after identity merging. A gender could not be identified automatically for the remaining unique authors in the dataset (about 25 M of them), which were excluded from further analysis.

*Geo-localization.* In order to quantitatively compare the effects of COVID-19 on contribution to public code around the world, we need to associate commit authors in the dataset to a geographic position, or “geolocate” them. More specifically, due to the granularity of COVID-19 data, we aim at geolocating commit authors to specific world *countries*.

As for gender detection, the scale of our data rules out manual approaches like interviewing contributors to ask where they are from—we need an automated country detection approach. We based ours on the email domain, which is recorded as part of commit metadata in our dataset.

We look up the email domain of each author in two datasets: country code top-level domain (ccTLD) and university websites, both retrieved from Wikipedia/Wikidata. The ccTLD dataset<sup>2</sup> associates top-level Internet domains (e.g., `.fr`, `.it`, `.mx`) to specific countries (e.g., France, Italy, and Mexico, respectively). The university website dataset associates full website URLs (such as `https://www.telecom-paris.fr/`, `https://www.unibo.it`, or `https://www.unam.mx`) to universities around the world, and the country where the university is incorporated. To obtain this dataset, we queried the Wikidata Query Service with a SPARQL query<sup>3</sup> returning all known universities on Wikipedia and, for each of them, its website and country of incorporation (both are available as Wikidata properties on university entities).

To associate authors to countries, we extract their location from their ccTLD and also match their email domains to university websites (stripping common and non-meaningful for us prefixes like URI schemes and `www.`). There are very few cases where a university domain is associated with a country that is different from the one associated with the ccTLD; we discard the authors in those cases. In case of failure of both lookups, we leave the authors non-geolocated and discard them from further analysis.

Because we were interested in matching authors to companies to answer RQ2, we also used Wikidata to retrieve a list of companies, their websites, and their countries of incorporation. We do not use this dataset to geo-localize authors, because a consistency check between the declared location of the company and the time zone of the majority of commits by the author (following the same method as described below) showed that this geo-localization method was less reliable than the ccTLD-based and university-based methods. However, we used this dataset to filter out authors whose email domain was associated with a multinational corporation, because we could not reliably associate them to a single country, or authors for which the ccTLD and reported company location were inconsistent.

Finally, for the present study, we were only interested in authors who created at least one commit in 2019, the year before the pandemic (because we wish to measure the impact of COVID-19 on preexisting contributors). Excluding authors without any contribution in 2019, and authors which we could not geo-localize, we obtain **a final set of 626 456 unique human contributors to public code, for which both gender and geographic location at country granularity were identified**, out of which 112 242 were located via their university and 562 314 via their ccTLD.

We validated the accuracy of our geolocation mechanism, by relying on the latest available dataset from the (now-defunct) GHTorrent project [26]

<sup>2</sup> A human-readable version of the ccTLD dataset can be found at [https://en.wikipedia.org/wiki/Country\\_code\\_top-level\\_domain](https://en.wikipedia.org/wiki/Country_code_top-level_domain), accessed November 2023. Note that we decided to discard a few ccTLDs that we considered problematic because they are widely used outside the country that they represent: `.ai`, `.cc`, `.co`, `.io`, `.me`, `.sh`, `.tv`, `.vg`, `.yt` (used in so-called “domain hacks”). We also excluded `.gp` because it is associated with several ISO3 country codes.

<sup>3</sup> The query can be found in the replication package of this paper [11].

as ground truth. The `User` table in GHTorrent includes both an `email` and a `country_code` field; the latter being an ISO 3166-1 alpha-2 country code.

The country code is obtained by GHTorrent applying Named Entity Recognition techniques to the `location` attribute self-declared by GitHub users on their profiles. Although the possibility exists that users provide fictitious or inaccurate locations, we contend that this is not the case for the vast majority. To support this assertion, we examined the commit timestamps of all users with a non-empty `country_code`, verifying whether their time zone offset is compatible with the associated country’s time zone offset at commit time. Our findings show that more than 90% of the commits are compatible with the declared country, assuming a reasonable amount of anomalies (e.g., developers being temporarily on a trip in a different country). This figure gives us confidence that the vast majority of self-declared locations are correct, and hence that it is reasonable to rely on GHTorrent location as ground truth for validation. (Note that GHTorrent is only useful for us as validation, rather than as a source for geolocating authors, because it is both limited to GitHub and largely outdated, missing a large proportion of our authors.)

Subsequently, we applied our geo-localization mechanism to the set of developers that have country information in the GHTorrent dataset, comparing our results with the `country_code` there. Out of the initial 1 127 255 entries, the classification successfully identified a country for 108 318 (9.6%) authors. Among the localized entries, the accuracy was 0.88, with a weighted mean precision of 0.91 and a weighted mean F1 score of 0.88. In the top 29 countries by support (number of entries in the validation set), accounting for the 90th percentile, only four entries had F1 score  $< 0.85$ .

*Number of commits per week.* For each gendered and geolocated unique human author, we count the number of commits authored each week (from Monday to Sunday), distinguishing in particular commits during working hours (defined as 7.00 AM–6.59 PM Monday to Friday) and outside working hours. We include in our dataset all commits made between Monday, May 28th, 2018 and Sunday, May 29th, 2022, to include about four full years of contribution activity data (209 complete weeks), including more than a full year of data before the onset of COVID-19, and without using any data for which the archiving by Software Heritage was not guaranteed to be complete. Note that in the dataset descriptive statistics (Table 1), the minimum number of commits in 2019 is 0 despite only considering authors that contributed at least once in 2019, because the aggregation at week level may lead to commits performed on the last few days of a year to be counted in the following year.

*Number of active days per week.* In addition to the number of commits, we also count the number of active days per week, defined as the number of days in the week in which the author authored at least one commit. This measure is used as an alternative specification of the dependent variable when answering RQ1.

**Table 1** Descriptive statistics of main variables

Variable	Observations	Mean	Std. dev.	Min	Max
Commits per week (all)	130 670 400	0.592	10.53	0	36 554
Commits per week (women)	17 051 160	0.381	4.162	0	4634
Commits per week (men)	113 619 240	0.624	11.18	0	36 554
Active days per week (all)	130 670 400	0.150	0.663	0	7
Active days per week (women)	17 051 160	0.106	0.546	0	7
Active days per week (men)	113 619 240	0.156	0.678	0	7
COVID-19 deaths per week per 100 000 inhabitants (2020+)	78 402 240	1.702	2.215	-0.489	103.9
Social distancing index (2020+)	77 864 036	0.505	0.217	0	1

Note: This table shows the basic descriptive statistics of the main variables used in the empirical analysis. The weekly number of COVID-19 related deaths per 100 000 inhabitants includes negative values because some countries (particularly at the beginning of the pandemic when the methodology to count COVID-19 related deaths was not yet stabilized) corrected the number of reported deaths from one week to the other, leading to some weeks with reported negative deaths to compensate for over reporting in the past. This, however, represents a small amount of country-week observations.

### 3.2 COVID-19 and demographic data

*COVID-19 deaths.* To know about weekly COVID-19 attributed deaths, we used the dataset “Daily cases and deaths by date reported to WHO” from the World Health Organization (WHO) Coronavirus (COVID-19) Dashboard,<sup>4</sup> which aggregates COVID-19 data by country/territory and varies per day. We aggregate the data per week.

*Lockdown stringency.* We use data from the Oxford COVID-19 Government Response Tracker [28] to measure the restrictiveness of the government’s pandemic policy. Notably, we use an original stringency index which captures the strictness of lockdown-style policies. This index reports values between 0 (no measure implemented) and 100 (highest stringency of social distance measures) and varies across countries and day. We aggregate the data per week and re-scaled it, making it vary between 0 and 1.

*Population.* As general population data, we relied on the “Demographic Indicators” dataset,<sup>5</sup> elaborated by the UN Department of Economic and Social Affairs Population Division. This data includes the total population of each country as of January 1st of every year. We divided the number of deaths per week of each year by the population on the 1st of January of the same year.

Table 1 shows the number of observations and basic descriptive statistics for the variables that were used in the empirical analysis.

<sup>4</sup> <https://covid19.who.int/data>, accessed November 2023

<sup>5</sup> <https://population.un.org/wpp/Download/Standard/MostUsed/> accessed November 2023

## 4 Methodology

To assess the causal effect of an event, we would ideally need to compare individuals in a world where the event took place, with the same individuals at the same time in an alternative world where the event did not happen. This is, of course, impossible. To approach this ideal scenario, scientists instead often set up to compare two randomly selected groups that only differ to the extent that one group is exposed to the event (the treated group) and the other is not (the control group). If the assignment to a group is random and the groups are large enough then, on average, the characteristics of the two groups are equal and the differences between them can be considered to occur only because of the exposure to the event of interest. However, in many settings, researchers are interested in studying the causal effect of an event, the exposure to which cannot be randomized across groups. This is the case of the COVID-19 pandemic. In these cases, several econometric techniques can be used to determine causal effects, depending on the context and the data structure. See the introduction of the recent paper by Graf-Vlachy and Wagner [27] for a summary of the four main econometric techniques that can be used to determine causal effects from observational data (and for more examples of why causal inference is important in software engineering research).

In this paper, we use a variation of a difference in differences (DID) strategy [4], one of the four main techniques mentioned above, which has already found applications in software engineering research [20, 33, 47], to estimate the causal impact of COVID-19 on the contribution of women to public code. This strategy takes advantage of the rich panel dimension of the data, i.e., being able to follow every author throughout their full contribution history.

The basic idea of DID is to compare a treated group and a control group that don't necessarily share the same characteristics. It is enough to assume that, in the absence of treatment, the trends for the outcome variable for the two groups would be parallel. In this setting, one group is exposed to the event and the other is not, but the two groups are not constructed randomly and can hence have different characteristics. Since we can observe the difference in the outcome variable between the two groups *before* the event of interest, we can subtract the difference in the outcome variable between the two groups (before the event) from the difference between the two groups after the event and obtain an estimation of how the treated group *would have looked like* in the absence of the event (assuming parallel trends). Any deviation from this estimation is considered to be the effect of the exposure to the event.

In our case, though, we do not have a true control group that we can consider never exposed to the event, because at some point or another, the entire world has been exposed to the COVID-19 pandemic. Therefore, we exploit the fact that different countries were exposed to COVID-19 at different points in time and with different intensities. In the case where there are several groups and several points in time, this method is called *two-way fixed effects (TWFE)* [4].

Finally, another difference in our analysis with respect to a classical DID strategy is that we are interested in the differential effect of COVID-19 exposure specifically for women. To capture this effect, we include the interaction between being exposed to the COVID-19 pandemic and being an author that we identify as a woman.

In the end, the main equation that we estimate is the following:

$$\begin{aligned} \text{Commits}_{acw} = & \beta_1 \text{Woman}_a \times \text{Exposure to COVID-19}_{cw} \\ & + \beta_2 \text{Exposure to COVID-19}_{cw} \\ & + \gamma_a + \theta_w + \alpha \text{Time} \times \text{Woman}_a + \epsilon_{acw} \end{aligned} \quad (1)$$

where:

- $\text{Commits}_{acw}$  is the number of commits of author  $a$  identified as living in country  $c$  during week  $w$ ;
- $\text{Woman}_a$  is a binary variable equal to 1 if author  $a$  is identified as being a woman or 0 otherwise;
- $\text{Exposure to COVID-19}_{cw}$  is a measure of exposure to the COVID-19 pandemic in country  $c$  during week  $w$ . We consider two alternative measures: the number of COVID-19 attributed deaths and the value of the social distance stringency index (see Section 3.2 for the origin of these data);
- $\gamma_a$  are author fixed effects;
- $\theta_w$  are week fixed effects;
- $\text{Time} \times \text{Woman}_a$  are women specific trends; and
- $\epsilon_{acw}$  is the error term.

Author fixed effects capture all author-constant characteristics that could explain the results in part, including, but not limited to, those related to being a woman in a specific country.

Imagine, for example, that in developing countries, authors create fewer commits than in developed countries and that developing countries are affected by COVID-19 more than developed countries. In this case, there would be a correlation between COVID-19 and the number of commits that is explained by different intrinsic characteristics of countries. However, in the absence of fixed effects, our estimates of a naive regression would capture this correlation, and we could misinterpret a negative coefficient of COVID-19 on the number of commits as causal, which would not be the case. The fixed effects that we include are not just at the country level but at the individual level, which allows us to get rid of any author-constant characteristics that could be correlated with COVID-19 (e.g., author-cohort, place of residence, or even gender).

More precisely, fixed effects subtract the mean of the fixed effects category. For author fixed effects, the mean of the number of commits per week of each author is subtracted from each observation, thus making the authors more comparable between each other.

Intuitively, a common way to interpret fixed effects is to imagine that the variation that is being used is within author, so that we compare the same author at different points in time.

Week fixed effects capture all time-related changes that affect all authors in the same way, including global trends. This could include for example the common effects of the WHO (World Health Organization) declaring a global pandemic for the whole planet, or a common decrease in economic activity across all countries. Other non-pandemic specific changes, for example a lower number of commits during summer or the Christmas season would also be captured by week fixed effects.

Finally, the women-specific trends capture time trends that are different for women. Including women-specific trends is important in this case because during the time period we consider for our analysis, we observe both an increasing slope for the total number of commits done by women and the number of COVID-19 deaths. Not controlling for a women-specific trend could result in misleading results, as we would correlate more COVID-19 deaths with more women contributions just because of a pre-existing trend.

The estimate of interest is  $\beta_1$  that captures the differential effect for women (with respect to comparable men) of being exposed to a certain level of COVID-19 in their country.

One assumption of classical ordinary least squares (OLS) regression is that all data points are independent of each other and are drawn from the same theoretical distribution. In many settings, this is not the case. Inhabitants of the same country, for example, share common institutions and characteristics that make them not independent of each other. Similarly, individuals of the same gender or age also share common characteristics. When this happens, some statistical methods allow relaxing the assumption that all observations are independent and identically distributed and determining some groups within which observations are not considered independent. This is called clustering.

Failing to recognize certain data points as not independent can lead to inaccurate standard errors and then to consider some coefficient estimates as statistically significant while in reality they are not. Clustering standard errors is then crucial for statistical inference. However, it is less clear at which level standard errors should be clustered. Recent econometric literature shows that for standard errors to be consistent<sup>6</sup> observations should be allowed to be correlated between units that share the same treatment.

Following the recent literature on standard errors [1] and DID estimators [7], we cluster standard errors at the treatment level. More precisely, we construct groups within which we allow correlation between observations. Data points of the same country, gender, and period of treatment (before or after COVID-19) are in the same group. Then, for each country, we have four different groups: one for women before the first week when any lockdown measure was implemented (or any COVID-19 related death had occurred); one for women after the week when the first lockdown measure was implemented in that country; and the same two groups for men.

---

<sup>6</sup> A *consistent estimator* is an estimator that, when the number of observations goes to infinity, converges to the real value of the parameter for the entire population.

To answer RQ1, we run the regressions corresponding to the empirical identification of equation (1) and cluster standard errors as described above using OLS with the statistical program STATA-17. To answer RQ2 and analyze which categories of women are most affected, we run the same regression as before, but restricting the sample to different categories. In the next section we describe the main result and detail further the analysis for different categories, showing obtained results.

## 5 Results

In this section, we present the results of our empirical analysis. Overall, the results show that higher exposure to lockdown measures due to COVID-19 reduced disproportionately the ability of women to contribute (compared to men). We also show that, among women, contributors that tended to contribute more outside working hours and contributors outside academia and companies are the most affected.

### 5.1 RQ1: differential effect of COVID-19 on women

The main result of our analysis, the overall (relative) effect of COVID-19 on women’s ability to contribute to public code, is shown in Table 2.

We consider two different measures of exposure to COVID-19: (1) the number of COVID-19-related deaths per 100 000 inhabitants in a given week in a given country, and (2) the stringency of the lockdown measures implemented as a response to the COVID-19 pandemic in a given week in a given country (captured by the index of lockdown-style policies strictness described in Section 3.2). We also consider two different measures of the ability of an author to contribute: the number of commits authored in a week, and the number of active days in a week (i.e., the number of days in which the author authored at least one commit).

With both measures of COVID-19 exposure and both measures of ability to contribute, we obtain results leading to the same conclusion: **COVID-19 disproportionately affected women’s ability to contribute to public code, compared to men.**

More precisely, in three out of four regressions in Table 2, the coefficients of interest (the interaction between the  $Woman_a$  binary variable and the two measures of COVID-19) are negative and statistically significant (at the 5% level in column 1 and at the 1% level in columns 2 and 4, which are the two columns reporting the results using the stringency index). For the last column (effect of the number of COVID-19 deaths on the number of active days), the coefficient is negative but not statistically significant.

Note that the magnitudes of our coefficients of interest do not represent the total effect of COVID-19 on women, but the *difference* of the effect between men and women. Furthermore, because of author fixed effects, they do not

**Table 2** Relative effect of COVID-19 on women’s contributing activity, using two different measures of COVID-19 and two different measures of contributing activity

	(1)	(2)	(3)	(4)
	N. of commits		N. of active days	
<b>Woman</b> × <b>COVID-19 deaths</b> (per 100 000 inhab.)	-0.00508** (0.00245)		-0.000589 (0.000423)	
<b>Woman</b> × <b>social distance index</b>		-0.0882*** (0.0240)		-0.0136*** (0.00499)
<b>Author FE</b>	Yes	Yes	Yes	Yes
<b>Week FE</b>	Yes	Yes	Yes	Yes
<b>Gender trend</b>	Yes	Yes	Yes	Yes
<b>Avg. dep. var.</b>	0.592	0.589	0.150	0.149
<b>Avg. (women)</b>	0.381	0.380	0.106	0.106
<b>Avg. (men)</b>	0.624	0.620	0.156	0.156
<b>Observations</b>	1.310e+08	1.300e+08	1.310e+08	1.300e+08
<b>N. clusters</b>	762	640	762	640

Note: This table shows the relative effect of COVID-19 on women’s contributing activity. Each column (numbered at the top for easier reference) corresponds to the results of a different regression. Columns 1 and 2 use number of commits as an outcome (dependent or Y variable), while columns 3 and 4 use the number of active days. Columns 1 and 3 use the number of COVID-19 deaths relative to the population as independent (or X) variable, while columns 2 and 4 use the social distance index. All columns include author fixed effects, week fixed effect, and worldwide gender-specific trends as indicated by the mention “Yes” in the corresponding line. The sample is all the authors that created at least one commit during 2019 that were localized in a country and that were assigned a gender. Woman is a binary variable indicating if the author is classified as woman. COVID-19 deaths (per 100 000 inhabitants) is the number of deaths due to COVID-19 per 100 000 inhabitants during a specific week in a specific country. The social distance index is a measure of the severity of the COVID-19 related social distance measures applied in each country each week. “Avg. dep. var.” is the average number of the dependent variable, i.e., number of commits per week for columns 1 and 2 or number of active days per week for columns 3 and 4, for our entire sample. To help interpret the results, we also provide the average number of the dependent variable for women authors and men authors in our sample. Standard errors are clustered at the gender-country-post treatment level. Stars represent different levels of statistical significance: \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.1$ .

allow comparing an average man with an average woman. Instead, we have to compare men and women authors that would have similar pre-pandemic contribution levels.

For column 1, the magnitudes of the estimates can be interpreted as follows. Consider a man and a woman that would have produced the same number of commits in a given week if no COVID-19 deaths had occurred in their country. Then, if instead one COVID-19 death per 100 000 inhabitants occurred, the woman would have produced 0.005 fewer commits than the man in this week. This relative decrease represents around 1.3% of the average number of per week contributions of women (computed by dividing the previous value by the

mean of the dependent variable for women only). Alternatively, one COVID-19 death per 100 000 inhabitants decreases the contribution of women (with respect to men) by around 2.1% of the difference in the averages between men and women, widening the gap of contributions.

To give a better sense of scale, we can interpret the magnitudes evaluating the effect for a specific number of COVID-19 deaths. During the week with the most number of COVID-19 deaths in the USA, there were about 7 deaths per 100 000 inhabitants, which then represents a 9.3% differential drop in women’s contributions compared to the mean of women and around a 15% increase in the average gap of contributions between men and women.

For column 2, magnitudes can be interpreted as follows. Changing the social distance index measure from no restrictions to the highest level of stringency decreases the contribution of women (again, with respect to comparable men) by 0.088 commits per week. This decrease represents around 23% of the average per-week contribution of women and around 37% of the average gap in the contributions between men and women.

Only 13 out of 216 countries reached the highest level of social distance stringency index. To give a better sense of the scale then, we can interpret the magnitudes for a change from no restrictions to a value of 0.75 (instead of 1) of the social distance stringency index, which was reached in 86% of countries at some point in time. The corresponding decrease in women’s contributions (relative to comparable men) represents 17% of the mean of women’s contributions and 28% of the gap between men and women.

For column 4, the magnitudes can be interpreted as follows. Changing the social distance index measure from no restrictions to the highest level of stringency decreases the number of active days of women (again, with respect to comparable men) by 0.0136 days per week. This decrease represents around 13% of the average number of active days for women contributors and around 27% of the average gap in the number of active days between men and women.

Although the two measures that we use to assess the impact of COVID-19 and the two measures that we use to evaluate the authors’ ability to contribute give similar results, in order to perform finer-grained analyses to answer RQ2, we focus on a main specification, which we will use on different subsamples of our main dataset (different categories of authors).

We choose to use social distancing and number of commits as our main specification. The social distance index is easier to interpret (as it varies from 0 to 1), and also better captures what is most likely to explain why COVID-19 had an impact on women’s ability to contribute to public code: the lockdown measures and their consequences in terms of the time that women (and men) could spend on such type of activity. Furthermore, this is the specification for which we obtain the lower p-value and the larger coefficient of interest (relative to the mean of the dependent variable). This makes more likely that we will have more statistical power (i.e., we will have a higher ability to capture existing effects) when reducing the number of observations to study the different subsamples.

### 5.1.1 Robustness checks

Table 7 in the appendix shows the robustness checks that we have conducted. They are variations of our main specification (whose results are repeated in column 1), either by altering data pre-processing, or by changing the regression specification itself.

In columns 2 and 3, we assess the robustness of our geo-localization method by changing some of its parameters: column 2 includes more authors, by lifting the exclusions of 9 ccTLDs that we considered problematic for geo-localization and of emails from multinational companies; on the opposite, column 3 uses a much more restricted list of ccTLDs, by excluding any ccTLD appearing on the Wikipedia page “Country code top-level domains with commercial licenses”<sup>7</sup> (in addition to our initial list of excluded ccTLDs), leading to a list of 55 exclusions, including large countries such as Iran and Italy.

In columns 4 to 6, we exclude some specific countries from our analysis to ensure that these countries are not exclusively driving our results. column 4 excludes the USA, which is both a large country with many contributors and one in which we could mostly identify contributors via university email domains (as the .us ccTLD is seldomly used); column 5 excludes countries which are outliers in terms of how many authors we identified relatively to their population (excluding both the top and bottom 5% countries when sorting by number of authors localized in this country over the country population). In particular, excluding the top 5% is another way of making sure that we are not sensible to countries whose ccTLD would be abused for domain hacks (and where we would localize too many authors compared to the reality). Finally, column 6 excludes countries for which we have too few women or men authors (less than 100), and which would therefore risk to lead to a biased comparison given the small number of observations.

In column 7, we test an even more demanding variant of our main specification, where gender-specific trends can differ in each different country. This alternative specification is likely to capture more gender-specific pre-pandemic variability due to varying politics across the world (e.g., while gender equality is slowly increasing in most countries, there are countries that experience setbacks).

The outcome for all these checks (shown in Table 7) **confirm the robustness of our main specification: the result continues to hold in all cases, with the same level of statistical significance (at the 1% level) and comparable coefficients.**

## 5.2 RQ2: who are the women most affected by COVID-19

To answer RQ2 we now re-run the regression for our main specification (social distancing index and number of weekly commits) on subsamples of our

<sup>7</sup> [https://en.wikipedia.org/wiki/Country\\_code\\_top-level\\_domains\\_with\\_commercial\\_licenses](https://en.wikipedia.org/wiki/Country_code_top-level_domains_with_commercial_licenses), accessed November 2023

**Table 3** Relative effect of COVID-19 on women’s number of weekly commits for different user categories, defined by year of first commit

Year of first commit	(1)	(2)	(3)	(4)
	all years	2019	2017–2018	before 2017
	<b>Number of commits</b>			
<b>Woman</b> × <b>social distance index</b>	-0.0882*** (0.0240)	-0.0513** (0.0249)	-0.0357 (0.0331)	-0.206*** (0.0489)
<b>Author FE</b>	Yes	Yes	Yes	Yes
<b>Week FE</b>	Yes	Yes	Yes	Yes
<b>Gender trend</b>	Yes	Yes	Yes	Yes
<b>Avg. dep. var.</b>	0.589	0.237	0.700	1.280
<b>Avg. (women)</b>	0.380	0.199	0.610	0.839
<b>Avg. (men)</b>	0.620	0.244	0.712	1.316
<b>Observations</b>	1.300e+08	6.680e+07	3.450e+07	2.840e+07
<b>N. clusters</b>	640	623	559	515

Note: This table shows the relative effect of COVID-19 on women’s contributing activity in different user categories depending on the year of first commit. Each column (numbered at the top for easier reference) corresponds to the results of a different regression. The sample is all the authors that created at least one commit during 2019 that were localized in a country and that were assigned a gender. Woman is a binary variable indicating if the author is classified as woman. The social distance index is a measure of the severity of the COVID-19 related social distance measures applied in each country each week. All columns include author fixed effects, week fixed effect, and worldwide gender-specific trends, as indicated by a “Yes” in the corresponding column and line. Column 1 is our baseline specification for comparison. Column 2-4 restrict the sample to specific categories depending on the year of their first commit. Column 2 shows the results for authors that wrote their first commit in 2019; column 3, in 2017 or 2018; and column 4, in 2016 or before. “Avg. dep. var.” is the average number of weekly commits for each subsample. To help interpret the results, we also provide the average number of weekly commits for women authors and men authors in each subsample. Standard errors are clustered at the gender-country-post treatment level. Stars represent different levels of statistical significance: \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.1$ .

population corresponding to various categories of contributors, in order to understand which women were the most affected by the COVID-19 pandemic (relatively to comparable men).

### 5.2.1 Distinguishing contributors by experience

We first categorize by contributor’s experience. For each author, we consider the year in which they created their first commit, and we create three categories out of this: started contributing in 2019 (the last year just prior to the pandemic), in 2017–2018, and before 2017. We run the same regression for each category. The results are presented in Table 3. Column 1 just repeats the main result from Table 2. The first category (column 2) is authors who created their first commit to public code in 2019. They represent  $\approx 51\%$  of our dataset. The second category (column 3) is authors who created their first commit in 2017–2018. They account for  $\approx 27\%$  of the dataset. Our third category (column 4)

is authors who created their first commit before 2017. They represent  $\approx 22\%$  of our dataset. By just looking at the mean of the dependent variable (overall, and restricted to specific genders), we can observe that the more experienced a contributor is, the more commits they contribute every week on average. (Note that our sample excludes authors that did not contribute any commit in 2019.)

The coefficients of interest are negative for all our categories of authors, and they are statistically significant for the first and last categories (authors with the least and the most experience).

Interpreting the magnitudes for the two categories with statistically significant coefficients, we can see that changing the social distance index from no restriction to the maximum level of restriction is responsible for a relative decrease of 0.206 weekly commits for the contributors with the most experience (column 4), compared to men that would have committed the same in a scenario without restrictions. This decrease corresponds to 25% of the average of weekly contributions for women in this category of more experienced users (compared to a 23% decrease in the main result), and a 43% increase in the gap between the average men and women levels of contribution (compared to a 37% increase of the gap in the main result).

For contributors with the least experience (column 2), the relative decrease of weekly commits is of 0.0513, which represents 26% of the average weekly contributions for women in this category, and a 114% increase of the gap between the average men and women levels of contribution in this category.

Overall, we can see that COVID-19 affected contributors at different levels of experience, at a similar level of magnitude (relative to the average number of weekly contributions by women in each category), and **we cannot conclude that a specific category of women, by contributing experience, is driving our main result.**

### 5.2.2 Distinguishing contributors by pre-pandemic levels of activity

We create a second categorization based on pre-pandemic contribution levels. For each author, we consider the number of commits they created in 2019 (before the onset of the pandemic) and we create three categories out of this. We run the same regression for each category. The results are presented in Table 4. Column 1 just repeats the main result from Table 2. The first category (column 2) is authors who created fewer than 100 commits in 2019; they represent  $\approx 91\%$  of the dataset. The second category (column 3) is authors who created between 100 and 999 commits in 2019;  $\approx 9\%$  of the dataset. The third category (column 4) is authors who created at least 1000 commits in 2019;  $\approx 0.4\%$  of the dataset. We keep statistically significant results for all the categories.

Interpreting the magnitudes, we can see that, for the first category (authors who committed fewer than 100 commits in 2019), the relative decrease of women contributions in a week with the maximal level of social distancing

**Table 4** Relative effect of COVID-19 on women’s number of weekly commits for different user categories, defined by number of commits in 2019

Number of commits in 2019	(1)	(2)	(3)	(4)
	all values	< 100	100–999	> 999
	<b>Number of commits</b>			
<b>Woman</b> × <b>social distance index</b>	-0.0882*** (0.0240)	-0.0604*** (0.0166)	-0.511*** (0.139)	-9.452*** (2.953)
<b>Author FE</b>	Yes	Yes	Yes	Yes
<b>Week FE</b>	Yes	Yes	Yes	Yes
<b>Gender trend</b>	Yes	Yes	Yes	Yes
<b>Avg. dep. var.</b>	0.589	0.249	2.993	28.83
<b>Avg. (women)</b>	0.380	0.202	2.357	20.08
<b>Avg. (men)</b>	0.620	0.257	3.065	29.40
<b>Observations</b>	1.300e+08	1.180e+08	1.120e+07	468300
<b>N. clusters</b>	640	638	460	200

Note: This table shows the relative effect of COVID-19 on women’s contributing activity in different user categories depending on the number of commits in 2019. Each column (numbered at the top for easier reference) corresponds to the results of a different regression. The sample is all the authors that created at least one commit during 2019 that were localized in a country and that were assigned a gender. Woman is a binary variable indicating if the author is classified as woman. The social distance index is a measure of the severity of the COVID-19 related social distance measures applied in each country each week. All columns include author fixed effects, week fixed effect, and worldwide gender-specific trends, as indicated by a “Yes” in the corresponding line and column. Column 1 is our baseline specification for comparison. Columns 2–4 restrict the sample to specific categories depending on the total number of commits of the author in 2019 (before the onset of the pandemic). Column 2 shows the results for authors that wrote less than 100 commits in 2019; column 3, between 100 and 999; and column 4, 1000 commits or more. “Avg. dep. var.” is the average number of weekly commits for each subsample. To help interpret the results, we also provide the average number of weekly commits for women authors and men authors in each subsample. Standard errors are clustered at the gender-country-post treatment level. Stars represent different levels of statistical significance: \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.1$ .

is of 0.0604 weekly commits, which represents 30% of the average weekly contribution for women in this category and a 110% increase of the gap between the average men and women contributions. For the second category (authors who committed between 100 and 999 commits in 2019), the relative decrease of women contributions in such a week is of 0.511 commits, which represents 22% of the average weekly contribution for women in this category and a 72% increase of the gap between the average men and women contributions. For the third category (authors who committed at least 1000 commits in 2019), the relative decrease of women contributions in such a week is of 9.452 commits, which represents 47% of the average weekly contribution for women in this category and a 101% increase of the gap between the average men and women contributions.

Overall, we can see that COVID-19 affected contributors with different levels of pre-pandemic activity and, similarly to the previous categorization, we

**Table 5** Relative effect of COVID-19 on women’s number of weekly commits for different user categories, defined by percentage of commits during working hours

% during working hours	(1) all values	(2) 0%	(3) ]0%,50%]	(4) ]50%,100%[	(5) 100%
<b>Number of commits</b>					
<b>Woman</b> × <b>social dist. index</b>	-0.0882*** (0.0240)	-0.0461*** (0.0177)	-0.170*** (0.0507)	-0.137*** (0.0460)	-0.00139 (0.0247)
<b>Author FE</b>	Yes	Yes	Yes	Yes	Yes
<b>Week FE</b>	Yes	Yes	Yes	Yes	Yes
<b>Gender trend</b>	Yes	Yes	Yes	Yes	Yes
<b>Avg. dep. var.</b>	0.589	0.0913	0.787	1.071	0.130
<b>Avg. (women)</b>	0.380	0.0698	0.515	0.720	0.117
<b>Avg. (men)</b>	0.620	0.0950	0.825	1.120	0.132
<b>Observations</b>	1.300e+08	1.770e+07	3.640e+07	3.860e+07	3.710e+07
<b>N. clusters</b>	640	534	559	571	596

Note: This table shows the relative effect of COVID-19 on women’s contributing activity in different user categories depending on the percentage of commits during working hours before 2020. Each column (numbered at the top for easier reference) corresponds to the results of a different regression. The sample is all the authors that created at least one commit during 2019 that were localized in a country and that were assigned a gender. *Woman* is a binary variable indicating if the author is classified as woman. The *social distance index* is a measure of the severity of the COVID-19 related social distance measures applied in each country each week. All columns include author fixed effects, week fixed effect, and worldwide gender-specific trends, as indicated by a “Yes” in the corresponding line and column. Column 1 is our baseline specification for comparison. Column 2–5 restrict the sample to specific categories, depending on the percentage of commits that each author wrote during working hours before 2020 (i.e., before the onset of the pandemic). Column 2 shows the results for authors that did not write any commits during working hours; column 3, for authors that wrote at least one commit during working hours but less than 50% of their commits are during working hours. Column 4 shows the results for authors that wrote the majority of their commits but not all during working hours; and column 5, for authors that wrote all of their commits during working hours. “Avg. dep. var.” is the average number of weekly commits for each subsample. To help interpret the results, we also provide the average number of weekly commits for women authors and men authors in each subsample. Standard errors are clustered at the gender-country-post treatment level. Stars represent different levels of statistical significance: \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.1$ .

**cannot conclude that a specific category of women, by pre-pandemic levels of activity, is driving our main result.**

### 5.2.3 Distinguishing professional and hobbyist contributors by contribution patterns

We create a third categorization to distinguish between professional and hobbyist contributors. For each author, we consider the percentage of commits they created during working hours in our commit dataset before 2020 (i.e.,

between May 28th, 2018 and the end of 2019). We define working hours as 7.00AM–6.59PM from Monday to Friday in commits’ local time. We create four categories out of this. We run the same regression for each category. The results are presented in Table 5. Column 1 just repeats the main result from Table 2. The first category (column 2) is authors who created all of their commits before 2020 outside working hours. They represent  $\approx 14\%$  of the dataset. The second category (column 3) is authors who created more than half of their commits before 2020 outside working hours;  $\approx 28\%$  of the dataset. The third category (column 4) is authors who created more than half of their commits before 2020 during working hours;  $\approx 30\%$  of the dataset. The last category (column 5) is authors who created all of their commits before 2020 during working hours;  $\approx 28\%$  of the dataset.

We obtain statistically significant results on a relative decrease of contributions of women due to COVID-19 for all categories except the last one (100% of commits before 2020 during working hours). Looking at the magnitudes allows us to understand a bit more which categories are most affected. Indeed, just by comparing the middle two categories (columns 3 and 4), we can already observe that the coefficient of interest is higher in absolute value for the category of authors that commit more outside working hours even though the average number of weekly commits for this category is lower.

More specifically, looking at the magnitudes for the three categories for which we got statistically significant results, we observe that:

- For the category of authors that create more than half of their commits during working hours (column 4), the relative decrease of weekly commits by women for a week with the highest levels of social distance measures is 0.137, which represents 19% of the average number of weekly commits for women in this category, and an increase of 34% of the gap between the average number of weekly commits of men and of women in this category.
- For the category of authors that create more than half of their commits outside working hours (column 3), the relative decrease of weekly commits by women for such a week is 0.170, which represents 33% of the average number of weekly commits for women in this category, and 55% of the gap between the average number of weekly commits of men and of women in this category.
- For the category of authors that create all their commits outside working hours (column 2), the relative decrease of weekly commits by women for such a week is 0.0461, which represents 66% of the average number of weekly commits for women in this category, and 183% of the gap between the average number of weekly commits of men and of women in this category.

A clear pattern emerges: **the more authors were contributing to public code outside of working hours before the pandemic, the more women were negatively impacted by the pandemic in their ability to contribute** (relatively to men in the same category and with the same pre-pandemic contribution levels).

**Table 6** Relative effect of COVID-19 on women’s number of weekly commits for different user categories, defined by their use of an email address from a university or a company

Email domain	(1)	(2)	(3)	(4)
	All	University	Company	Non uni. & non company
	<b>Number of commits</b>			
<b>Woman</b> × <b>social dist. index</b>	-0.0882*** (0.0240)	-0.0492 (0.0550)	-0.0785 (0.134)	-0.104*** (0.0255)
<b>Author FE</b>	Yes	Yes	Yes	Yes
<b>Week FE</b>	Yes	Yes	Yes	Yes
<b>Gender trend</b>	Yes	Yes	Yes	Yes
<b>Avg. dep. var.</b>	0.589	0.602	0.504	0.589
<b>Avg. (women)</b>	0.380	0.425	0.421	0.366
<b>Avg. (men)</b>	0.620	0.642	0.513	0.619
<b>Observations</b>	1.300e+08	2.350e+07	2.370e+06	1.040e+08
<b>N. clusters</b>	640	438	272	632

Note: This table shows the relative effect of COVID-19 on women’s contributing activity in different user categories depending on whether they have a company or university email address. Each column (numbered at the top for easier reference) corresponds to the results of a different regression. The sample is all the authors that created at least one commit during 2019 that were localized in a country and that were assigned a gender. Woman is a binary variable indicating if the author is classified as woman. The social distance index is a measure of the severity of the COVID-19 related social distance measures applied in each country each week. All columns include author fixed effects, week fixed effect, and worldwide gender-specific trends, as indicated by a “Yes” in the corresponding line and column. Column 1 is our baseline specification for comparison. Column 2–4 restrict the sample to specific categories, depending on whether they were identified as belonging to a university, to a company or none of these based on their email address. Column 2 shows the results for authors with a university email address, column 3 shows the results for authors with a company email address, while column 4 shows the results for the authors that were not identified as belonging to a university or a company. “Avg. dep. var.” is the average number of weekly commits for each subsample. To help interpret the results, we also provide the average number of weekly commits for women authors and men authors in each subsample. Standard errors are clustered at the gender-country-post treatment level. Stars represent different levels of statistical significance: \*\*\* p<0.01; \*\* p<0.05; \* p<0.1.

#### 5.2.4 Distinguishing professional and hobbyist contributors by their use of a professional email address

We create a fourth categorization to distinguish between professional and hobbyist contributors, this time based on their use of a professional email address. For each author, we consider whether we were able to match their email address to a university or a company, using data from Wikidata. We create three categories, and we run the same regression for each category. The results are presented in Table 6. Column 1 just repeats the main result from Table 2. The first category (column 2) is authors that use a university-based email address that we could recognize. They represent  $\approx 17\%$  of our dataset. The second

category (column 3) is authors that use a company-based email address that we could recognize. They represent  $\approx 1.7\%$  of our dataset. The last category (column 4) is authors that use an email address that we could not recognize as being from a university or a company.

We only get statistically significant results for the last category (column 4). The magnitude of the coefficient of interest is higher for this category compared to our main result, even though the average number of weekly commits in this category is slightly smaller. From this, we conclude that our **main result is mostly driven by contributors outside universities and companies**, i.e., mostly driven by hobbyists. Note that by excluding university email addresses, we exclude both professional academics and students using their university email address. For most universities, it is impossible to distinguish between students and employees based on the email address alone.

Interpreting the magnitude for this last category, we observe that the relative decrease of women contributions for weeks with the highest levels of social distance measures is of 0.104 weekly commits, which represents 28% of the average weekly commit number by women in this category and 41% of the gap between the averages for men and women in this category.

## 6 Discussion

In this section we discuss our empirical findings (from Section 5), providing plausible interpretations as of their origin, and evaluating threats to their validity.

### 6.1 Interpretation of the findings

Our main result (RQ1, Table 2) shows that **the COVID-19 pandemic induced (rather than being merely correlated with) an impairment in the ability of women to contribute to public code**, relatively to men. The result holds for different ways of measuring the ability to contribute, and is statistically strong, both in terms of significance and robustness (see Section 5.1.1).

While the goal of this work was primarily to verify *if* such a causal relationship existed and its magnitude, the breakdown by different groups of contributors that we conducted to answer RQ2 suggests interpretations as to where the impairment comes from.

On the one hand, our first and second categorizations (by experience and by pre-pandemic levels of activity) did not show a specific category of contributors being more affected than others. We still obtained statistically significant results showing women being disproportionately affected by the pandemic in virtually all of these groups. Looking into the magnitudes of coefficients, we observed some variations in the way certain categories were affected, but with no clear pattern suggesting that a specific category was more affected than others.

On the other hand, our third and fourth categorizations (by contribution patterns and by the use of a professional email address) revealed a clear difference in the way the pandemic affected hobbyist and professional contributors. Our two methods of distinguishing between these two types of contributors led to the same conclusion: **hobbyist contributors were more affected by the pandemic than professional contributors.**

In the case of the third categorization, the more women were contributing outside of working hours before the pandemic, the more affected they were, relatively to their previous contribution levels, and relatively to men in the same category. We also did not obtain any statistically significant differential effect for authors that were contributing all their commits during working hours before the pandemic.

In the case of the fourth categorization, contributors outside academia and companies were significantly affected, but for contributors inside academia and companies, we did not obtain any statistically significant result showing a differential effect of the pandemic on women's versus men's contributions. This does not mean that COVID-19 did not affect these contributors, but rather that the effect was not significantly different for men and women.

Overall, and consistently with the literature in other domains, one possible explanation for these results is that women were disproportionately affected by additional burdens because of social distancing measures (e.g., taking care of children when schools are closed), and therefore had less spare time (in comparison with men with similar pre-pandemic contribution patterns) that they could allocate, in particular, to contributing to open source as a hobby. On the other hand, women who contribute to public code as part of their work are not affected as much, because work during lockdown could (and did in many cases) continue.

In fact, a tradition of working from home has existed in software development since well before COVID-19. It is plausible that this pre-existing work-from-home culture is partially responsible for the reduced, or lack of differential effect of the pandemic that we observe for professional contributors. Developers that were *already* remote working have been impacted nonetheless by the pandemic [5, 10, 23, 35] (as discussed in detail in Section 2).

For the case of our negative result for contributors with university email addresses, we can speculate different explanations for students and university employees that are consistent with our results. For university employees, the explanation is most likely the same as for contributors from companies: they continued to work during the lockdown, and therefore their ability to contribute to public code was not affected by the pandemic (relatively to men). Regarding students, women students are less likely to be disproportionately affected (relatively to their men counterparts) by social distancing measures because most of them do not yet have children or elderly parents in their care. For some students, the pandemic could even have increased the spare time that they could allocate to contributing to public code, but there is no reason that this effect would be different for men and women students.

## 6.2 Threats to validity

In the remainder of this section we discuss threats to the validity of the reported findings and the mitigations put in place to counter them. We follow the terminology and threats classification by Runeson et al. [45].

### 6.2.1 Construct validity

*Analysis scale.* We have conducted a large-scale analysis that encompasses, after data cleanup and identity merging, a set of 41 M authors. At this scale, manual approaches (e.g., interviews) are not viable to obtain information about the gender or geographic location of the entire population. Joining identities from the initial dataset with external platforms (e.g., social media websites) was also not feasible due to the lack of a viable join key. We have hence resorted to fully automated methods for detecting gender and geographic location, in part reusing tools and techniques from recent works in the literature (for gender detection at this scale) and in part exploiting signals that we expect to be reliable (e.g., university email domains). We therefore do not consider that the chosen techniques for feature detection pose a significant risk to the validity of the obtained results. Nonetheless, we critically review and discuss below specific aspects of our choices for gender detection and geolocation.

*Gender detection.* We largely share the threats to the validity of gender detection of Rossi and Zacchiroli [44], which we briefly recall here. Several automated tools and services for name-based gender detection exist. Among them, we have chosen to build upon `GENDER-GUESSER` for two main reasons: (1) it is open source, which not only enables better replicability and verifiability, but also allows scaling to the number of authors in our dataset without incurring per-query costs; (2) according to the benchmark conducted by Santamaria et al. [48], `GENDER-GUESSER` performs comparatively well with diverse datasets containing authors from all over the world, like ours. Given `GENDER-GUESSER` works on *first* names, our methodology adds on top of it two levels of majority assessment for determining the gender of an author: one among the tokens of each of their identities, another among the multiple identities of the same author. One can easily construct artificial cases in which legitimate family names skew the gender of an author in the wrong direction, but this does not appear to happen in samples of the dataset that we have manually verified.

We were unable to determine the gender of significant parts of the dataset, but this is common for such large-scale experiments. Also, the remaining sample that we have analyzed is still very significant in size, especially w.r.t. most of the studies in the field.

*Geolocation.* We have used two different techniques for geolocation, both based on author emails: country inference based on country code top-level domains (ccTLDs) and inference based on university domains.

Both techniques can be gamed by authors, as emails in Git commits are not verified, but there is little incentive to do that just to be assigned to a different country. Domain hacks (e.g., `.io` ccTLD for technology-related domains or `.me` for personal websites) will result in misclassifications, which we have accounted for by excluding ccTLDs that are commonly used outside their respective countries. We also performed a robustness check, where we exclude an even larger sample of ccTLDs that could be misused. On the other hand, for historical reasons people in the US tend to use much less their own ccTLD domain (`.us`) than people located elsewhere in the world; such population in our experiment will be geolocated in the US mainly via the university domain technique, possibly resulting in their under-representation. For this reason, we also performed a robustness check where authors from the US were excluded.

We have verified the reliability of geolocation in two ways. First, we have computed the agreement between the two geolocation techniques and found it to be very high: within the population of authors having contributed at least one commit in 2019 for which we could detect the gender and that were geolocated with *both* techniques, only 40 unique authors were assigned to a different country. We have excluded these authors from the analysis.

Second, we have compared our geolocation results to the self-declared location of authors on GitHub, using as ground truth the latest available GHTorrent [26] data dump, obtaining high accuracy (see Section 3.1 for details). These results instill confidence in our geo-localization technique, in spite of the amount of unclassifiable entries. Notably, the majority of unclassified email addresses belong to familiar commercial domains such as Gmail, Hotmail, Yahoo, and anonymized GitHub addresses, totaling 66% (441 938 out of 671 766) of all distinct email domains in the validation dataset.

As with all classification methods that introduce data filtering, there is a risk that filtering introduces a bias with respect to the dimensions of interest for the study that is conducted. This seems very unlikely in this case: while the filtering can introduce a gender and/or a geographic bias to the data, we are interested in the temporal evolution of these data and the biases are unlikely to change through time.

### 6.2.2 External validity

*Public code coverage.* We started from a very large, but necessarily incomplete, sample of the entire body of public code and its version control history. No archive of public code can be complete, and we inherit the incompleteness of the archive we used as starting point: Software Heritage. To the best of our knowledge, Software Heritage is the largest publicly accessible archive of public code, which we believe validates our choice of it as the initial dataset.

Nonetheless, we acknowledge that this implies our analysis is affected by the archival coverage, or lack thereof, that Software Heritage has of public software development platforms. In terms of archived platforms, Software Heritage archives the most popular development forges (e.g., GitHub and GitLab.com),

as well as a number of less popular ones,<sup>8</sup> so we consider the risk of missing platforms entirely to be limited. But Software Heritage might also be missing relevant development activities (most importantly for our methodology: commits) due to various reasons. We have discussed this with the archive operators and learned that there is indeed an archival lag w.r.t. GitHub and we have taken it into account. Specifically, we have verified that archive operators consider the data dump we started from (March 18th, 2024) to be complete up to May 30th, 2022, and restricted our analysis to that date: which was more than enough for our needs, covering more than 2 years from the beginning of the pandemic (early 2020).

*Gender identification.* Name-based gender identification is not a perfect method and its application could result in some form of bias because of technical and social related issues. In the data processing described in Section 3.1, we were able to associate a gender only to a part of the authors, leaving several in the *unknown* class. Our observations about that *unknown* population support the idea that they mostly consist of a specific, well-characterized type of contributors, rather than a group of contributors that we randomly could not assign to a gender. Focusing on the May 2018–May 2022 period, while the unknown active authors account for 61.4% of the total, the commits they performed in that time span amount to 43.7% of the total; similarly, the average number of commits performed by unknown authors is 35.5, which is significantly lower than the 53.9 of women and 77.5 of men. One possible explanation for this difference is that these unknown are largely constituted by authors that did not configure a Git client with their name, e.g., because they are committing directly from a GitHub account. In fact, when looking at the authors that we are able to geolocate, there are fewer authors for which we could not detect the gender: 38% for authors localized with their email ccTLD and 40% for authors localized with their university email domain.

Due to systemic online discrimination and harassment, women can take steps to hide their gender in public interactions, including code contributions. They can for instance use a pseudonym or adopt a male-looking name. This will result in under-counting women authors and their contributions (if they adopt a pseudonym that will be detected as “unknown” by `GENDER-GUESSER`) and possibly also in over-counting men authors and their contributions (if they adopt a male name or a pseudonym that will be detected as “male”). No mitigation is possible for this in the realm of name-based gender detection. We only observe that we consider unlikely that this phenomenon has changed significantly across the pandemic and that the worst case (in terms of validity risk), women adopting male names, is less likely than other cases.

*Non geo-localizable developers.* The analysis we conducted is restricted to developers that we could geolocate with either of the email-based methods we

---

<sup>8</sup> See <https://archive.softwareheritage.org> for coverage details, accessed December 2023.

have used. It is possible that the resulting population subset is not representative of global tendencies, and also that country-specific patterns have been excluded from our analysis due to non-geolocated developers in those countries. This risk cannot be avoided entirely, hence we make no claims about the generality of our claims outside the analyzed population sample. Nonetheless, we conducted and reported about robustness checks (see Section 5.1.1) to verify this, not finding any issue.

### 6.2.3 Reliability

Reliability threats concern the extent to which our analyses and results depend on the team obtaining them. As it is common practice, we mitigate this risk by providing a replication package for the work presented in this paper [11].

Note that, due to the presence of personal information in the dataset (names and emails) we are unable to provide a *complete* replication package, in particular we have excluded such information from the replication package. As the next best alternative, we include in the replication package the list of commit identifiers (in the form of SWHIDs [18]) that constitute our starting public code dataset in Figure 1. This enables independent crawling of the same information, as well as cross-referencing the relevant commits with other datasets.

From there on, all other results are verifiable and reproducible using the replication package.

## 7 Conclusion and future work

To act in favor of more inclusion of women in free/open source software (FOSS) development, it is crucial to understand the causes of their current underrepresentation. In this work, we have investigated one such (recent and worldwide) cause, which had not been studied before in the context of FOSS development: the arrival of the COVID-19 pandemic.

By using a large dataset of contributors and contributions to public code, obtained from the Software Heritage archive, and difference in differences (DID), we have shown that there exists a causal relationship between the pandemic and a disproportionate reduction in women’s ability to contribute to public code, relative to men. Furthermore, we have looked into which groups of women contributors have been most impacted by the pandemic, finding that it is specifically among hobbyist contributors that women have been most disproportionately affected.

These results are consistent with those obtained for contexts in society other than computing and open source, showing how systemic household disparities have been heightened by the pandemic. We argue that open source communities who are committed to reducing gender inequalities need to take into account such systemic disparities and their exacerbation by external

shocks, like the COVID-19 pandemic, when designing, implementing, and monitoring policies aimed at increasing diversity and inclusion.

*Future work.* Several aspects of this phenomenon remain to be investigated as future work.

Several years have now passed since the beginning of the COVID-19 pandemic. Stringency measures to counter it have been weakened around the world, for better or worse. It is now possible to verify whether doing so has enacted a “bounce back” in the ability of women to contribute to public code, relative to men. While similar large-scale analyses should be possible, the methodology we have used in this work is not directly applicable to this new question, as it directly relates weeks with COVID-19 stringency measures to the contribution activity of authors in the same week. Thus, this method is not suited to look into longer-term effects of the pandemic on women’s ability to contribute to public code.

In terms of policy factors, it would be interesting to know if some open source communities were better able to face the COVID-19 impact on women participation than others. To verify that, the subgroup analysis that we conducted in this paper can be refined to partition contributors by the usual project characteristics (size, longevity, technological stack, etc.) as well as by relevant factors for the problem scope, e.g., the adoption of code of conducts (CoC) or other DEI (Diversity, Equity, Inclusion) related actions. This finer-grained quantitative analysis could be complemented by more qualitative studies, e.g., interviews or surveys with maintainers and contributors of projects that have been particularly successful in maintaining or increasing the level of women participation despite the pandemic, to understand what practices have been effective.

Finally, we would like to look into the nature of the work done by contributors in our dataset. Were women contributors working on specific tasks (e.g., development, testing, documenting, quality assurance, etc.) more impacted than women working on others? And, if so, why is that the case? Insights obtained on these questions can help make FOSS communities more resilient to future major external shocks, like COVID-19 has been.

## Conflicts of Interest

The authors declare that they have no conflict of interest.

## References

1. Alberto Abadie, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. When should you adjust standard errors for clustering? *The Quarterly Journal of Economics*, 138(1):1–35, 2023.
2. Jean-François Abramatic, Roberto Di Cosmo, and Stefano Zacchiroli. Building the universal archive of source code. *Communications of the ACM*, 61(10):29–31, October 2018.

3. Abi Adams-Prassl, Teodora Boneva, Marta Golin, and Christopher Rauh. Inequality in the impact of the coronavirus shock: Evidence from real time surveys. *Journal of Public Economics*, 189:104245, 2020.
4. Joshua D. Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics*. Princeton University Press, Princeton, NJ, USA, April 2009.
5. Lingfeng Bao, Tao Li, Xin Xia, Kaiyu Zhu, Hui Li, and Xiaohu Yang. How does working from home affect developer productivity? — A case study of Baidu during the COVID-19 pandemic. *Science China Information Sciences*, 65(4):142102, 2022.
6. Michele Belot, Syngjoo Choi, Julian C. Jamison, Nicholas W. Papageorge, Egon Tripodi, and Eline van den Broek-Altenburg. Six-Country Survey on Covid-19, May 2020. [Online; accessed 30. Nov. 2023].
7. Marianne Bertrand, Esther Dufo, and Sendhil Mullainathan. How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*, 119(1):249–275, 2004.
8. Richard Blundell, Monica Costa Dias, Robert Joyce, and Xiaowei Xu. COVID-19 and Inequalities. *Fiscal Studies*, 41(2):291–319, 2020.
9. Carmen Botella, Silvia Rueda, Emilia López-Iñesta, and Paula Marzal. Gender diversity in STEM disciplines: A multiple factor problem. *Entropy*, 21(1):30, 2019.
10. Jenna Butler and Sonia Jaffe. Challenges and Gratitude: A Diary Study of Software Engineers Working From Home During Covid-19 Pandemic. In *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, pages 362–363, 2021.
11. Annalí Casanueva, Davide Rossi, Stefano Zacchiroli, and Théo Zimmermann. The impact of the COVID-19 pandemic on women’s contribution to public code — replication package. <https://doi.org/10.5281/zenodo.12773108>, 2024.
12. Theophania Chavatzia. Cracking the code: Girls’ and women’s education in science, technology, engineering and mathematics (STEM). Technical report, UNESCO, 2017. Online; accessed 2022-02-08.
13. Caitlyn Collins, Liana Christin Landivar, Leah Ruppanner, and William J. Scarborough. COVID-19 and the gender gap in work hours. *Gender, Work & Organization*, 28(S1):101–112, 2021.
14. Paulo Anselmo da Mota Silveira Neto, Umme Ayda Mannan, Eduardo Santana de Almeida, Nachiappan Nagappan, David Lo, Pavneet Singh Kochhar, Cuiyun Gao, and Iftekhar Ahmed. A deep dive into the impact of COVID-19 on software development. *IEEE Trans. Software Eng.*, 48(9):3342–3360, 2022.
15. Hai-Anh H. Dang and Cuong Viet Nguyen. Gender inequality during the COVID-19 pandemic: Income, expenditure, savings, and job loss. *World Development*, 140:105296, 2021.
16. Daniela Del Boca, Noemi Oggero, Paola Profeta, and Mariacristina Rossi. Women’s and men’s work, housework and childcare, before and during COVID-19. *Review of Economics of the Household*, 18(4):1001–1017, 2020.
17. Tapajit Dey, Sara Mousavi, Eduardo Ponce, Tanner Fry, Bogdan Vasilescu, Anna Filippova, and Audris Mockus. Detecting and characterizing bots that commit code. In Sunghun Kim, Georgios Gousios, Sarah Nadi, and Joseph Hejderup, editors, *MSR ’20: 17th International Conference on Mining Software Repositories, Seoul, Republic of Korea, 29-30 June, 2020*, pages 209–219. ACM, 2020.
18. Roberto Di Cosmo, Morane Gruenpeter, and Stefano Zacchiroli. Identifiers for digital objects: the case of software source code preservation. In *Proceedings of the 15th International Conference on Digital Preservation, iPRES 2018, Boston, USA, September 2018*.
19. Naveen Donthu and Anders Gustafsson. Effects of COVID-19 on business and research. *Journal of Business Research*, 117:284–289, 2020.
20. Hongbo Fang, Hemank Lamba, James Herbsleb, and Bogdan Vasilescu. ”This is damn slick!”: estimating the impact of tweets on open source project popularity and new contributors. In *Proceedings of the 44th International Conference on Software Engineering*, pages 2116–2129, Pittsburgh Pennsylvania, May 2022. ACM.
21. Lidia Farre, Yarine Fawaz, Libertad Gonzalez, and Jennifer Graves. How the Covid-19 Lockdown Affected Gender Inequality in Paid and Unpaid Work in Spain, July 2020. [Online; accessed 30. Nov. 2023].

22. Israel Saeta Pérez Ferhat Elmas, David Arcos. gender-guesser. <https://github.com/lead-ratings/gender-guesser>, 2015. Retrieved 2021-09-28.
23. Denae Ford, Margaret-Anne Storey, Thomas Zimmermann, Christian Bird, Sonia Jaffe, Chandra Maddila, Jenna L. Butler, Brian Houck, and Nachiappan Nagappan. A Tale of Two Cities: Software Developers Working from Home during the COVID-19 Pandemic. *ACM Transactions on Software Engineering and Methodology*, 31(2):27:1–27:37, 2021.
24. Karen A Frenkel. Women and computing. *Communications of the ACM*, 33(11):34–46, 1990.
25. Taiyo Fukai, Masato Ikeda, Daiji Kawaguchi, and Shintaro Yamaguchi. COVID-19 and the employment gender gap in Japan. *Journal of the Japanese and International Economies*, 68:101256, June 2023.
26. Georgios Gousios and Diomidis Spinellis. Ghtorrent: Github's data from a firehose. In Michele Lanza, Massimiliano Di Penta, and Tao Xie, editors, *9th IEEE Working Conference of Mining Software Repositories, MSR*, pages 12–21. IEEE Computer Society, 2012.
27. Lorenz Graf-Vlachy and Stefan Wagner. Cleaning Up Confounding: Accounting for Endogeneity Using Instrumental Variables and Two-Stage Models. *ACM Transactions on Software Engineering and Methodology*, page 3674730, June 2024.
28. Thomas Hale, Samuel Webster, Anna Petherick, Toby Phillips, and Beatriz Kira. Oxford covid-19 government response tracker. *Blavatnik School of Government*, 25, 2020.
29. Catherine Hill, Christianne Corbett, and Andresse St Rose. *Why so few? Women in science, technology, engineering, and mathematics*. ERIC, 2010.
30. Meenakshi Kaushik and Neha Guleria. The Impact of Pandemic COVID -19 in Workplace. *European Journal of Business and Management*, 12(15):9–18, 2020.
31. Vanessa Klotzman, Farima Farmahinifarahani, and Cristina Lopes. Public Software Development Activity During the Pandemic. In *Proceedings of the 15th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, ESEM '21, pages 1–12. Association for Computing Machinery, 2021.
32. Letícia S. Machado, Clara Marques Caldeira, Marcelo Gattermann Perin, and Cleidson R. B. de Souza. Gendered experiences of software engineers during the COVID-19 crisis. *IEEE Softw.*, 38(2):38–44, 2021.
33. Danaja Maldeniya, Ceren Budak, Lionel P. Robert Jr., and Daniel M. Romero. Herding a Deluge of Good Samaritans: How GitHub Projects Respond to Increased Attention. In *Proceedings of The Web Conference 2020*, pages 2055–2065. Association for Computing Machinery, New York, NY, USA, April 2020.
34. Jane Margolis and Allan Fisher. *Unlocking the clubhouse: Women in computing*. MIT press, 2002.
35. Grant R. McDermott and Benjamin Hansen. Labor reallocation and remote work during COVID-19: Real-time evidence from GitHub. *NBER*, December 2021.
36. Nikola Pejić, Zaharije Radivojević, and Miloš Cvetavnović. Analyzing the Impact of COVID-19 on GitHub Event Trends. *Sustainability*, 15(19):14622, 2023.
37. Richard J. Petts, Daniel L. Carlson, and Joanna R. Pepin. A gendered pandemic: Childcare, homeschooling, and parents' employment during COVID-19. *Gender, Work & Organization*, 28(S2):515–534, 2021.
38. Antoine Pietri, Diomidis Spinellis, and Stefano Zacchiroli. The Software Heritage graph dataset: public software development under one roof. In Margaret-Anne D. Storey, Bram Adams, and Sonia Haiduc, editors, *Proceedings of the 16th International Conference on Mining Software Repositories, MSR 2019, 26-27 May 2019, Montreal, Canada.*, pages 138–142. IEEE / ACM, 2019.
39. Kate Power. The COVID-19 pandemic has increased the care burden of women and families. *Sustainability: Science, Practice and Policy*, 16(1):67–73, 2020.
40. Gede Artha Azriadi Prana, Denae Ford, Ayushi Rastogi, David Lo, Rahul Purandare, and Nachiappan Nagappan. Including Everyone, Everywhere: Understanding Opportunities and Challenges of Geographic Gender-Inclusion in OSS. *IEEE Transactions on Software Engineering*, 48(9):3394–3409, September 2022.
41. Paul Ralph, Sebastian Baltes, Gianisa Adisaputri, Richard Torkar, Vladimir Kovalenko, Marcos Kalinowski, Nicole Novielli, Shin Yoo, Xavier Devroey, Xin Tan, Minghui Zhou, Burak Turhan, Rashina Hoda, Hideaki Hata, Gregorio Robles, Amin Milani Fard, and

- Rana Alkadhi. Pandemic programming. *Empirical Software Engineering*, 25(6):4927–4961, 2020.
42. Malte Reichelt, Kinga Makovi, and Anahit Sargsyan. The impact of COVID-19 on gender inequality in the labor market and gender-role attitudes. *European Societies*, 23:S228–S245, 2021.
  43. Anni Reinking and Barbara Martin. The gender gap in STEM fields: Theories, movements, and ideas to engage girls in STEM. Technical report, University of Alicante, 2018.
  44. Davide Rossi and Stefano Zacchiroli. Worldwide gender differences in public code contributions and how they have been affected by the COVID-19 pandemic. In *44th IEEE/ACM International Conference on Software Engineering: Software Engineering in Society ICSE (SEIS) 2022, Pittsburgh, PA, USA, May 22-24, 2022*, pages 172–183. IEEE/ACM, 2022.
  45. Per Runeson and Martin Höst. Guidelines for conducting and reporting case study research in software engineering. *Empirical Software Engineering*, 14(2):131, 2009.
  46. Daniel Russo, Paul H. P. Hanel, Seraphina Altnickel, and Niels van Berkel. The daily life of software engineers during the COVID-19 pandemic. In *43rd IEEE/ACM International Conference on Software Engineering: Software Engineering in Practice, ICSE (SEIP) 2021, Madrid, Spain, May 25-28, 2021*, pages 364–373. IEEE, 2021.
  47. Giuseppe Russo Latona, Christoph Gote, Christian Zingg, Giona Casiraghi, Luca Verginer, and Frank Schweitzer. Shock! Quantifying the Impact of Core Developers’ Dropout on the Productivity of OSS Projects. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 706–709, Singapore Singapore, May 2024. ACM.
  48. Lucía Santamaría and Helena Mihaljevic. Comparison and benchmark of name-to-gender inference services. *PeerJ Computer Science*, 4:e156, 2018.
  49. Almudena Sevilla and Sarah Smith. Baby steps: The gender division of childcare during the COVID-19 pandemic. *Oxford Review of Economic Policy*, 36:S169–S186, 2020.
  50. Bianca Trinkenreich, Ricardo Britto, Marco A. Gerosa, and Igor Steinmacher. An empirical investigation on the challenges faced by women in the software industry: A case study. In *Proceedings of the 2022 ACM/IEEE 44th International Conference on Software Engineering: Software Engineering in Society, ICSE-SEIS ’22*, pages 24–35, New York, NY, USA, October 2022. Association for Computing Machinery.
  51. Bianca Trinkenreich, Marco Aurélio Gerosa, and Igor Steinmacher. Women in open source: We need to talk about it. *Computer*, 55(12):145–149, 2022.
  52. Bianca Trinkenreich, Igor Wiese, Anita Sarma, Marco Gerosa, and Igor Steinmacher. Women’s Participation in Open Source Software: A Survey of the Literature. *ACM Transactions on Software Engineering and Methodology*, 31(4):81:1–81:37, August 2022.
  53. Ming-Te Wang and Jessica L Degol. Gender gap in science, technology, engineering, and mathematics (STEM): Current knowledge, implications for practice, policy, and future directions. *Educational psychology review*, 29(1):119–140, 2017.
  54. Stefano Zacchiroli. Gender differences in public code contributions: a 50-year perspective. *IEEE Software*, 38(2):45–50, 2021.
  55. Gema Zamorro and María J. Prados. Gender differences in couples’ division of childcare, work and mental health during COVID-19. *Review of Economics of the Household*, 19(1):11–40, 2021.

## **A Robustness checks results**

Table 7 shows the outcomes of the robustness checks we have conducted on our main result. (See Section 5.1.1 for a discussion of these checks.)

Table 7 Robustness checks on the main specification: Relative effect of COVID-19 on women’s number of weekly commits

Variant	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	main	no exclusion	no commerc. ccTLD	excl. USA	excl. low repr.	excl. few obs.	country-gender trend
<b>Number of commits</b>							
<b>Woman</b> × <b>social distance index</b>	-0.0882*** (0.0240)	-0.0898*** (0.0234)	-0.0909*** (0.0267)	-0.0845*** (0.0252)	-0.0877*** (0.0240)	-0.0961*** (0.0236)	-0.0828*** (0.0209)
<b>Author FE</b>	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<b>Week FE</b>	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<b>Gender trend</b>	Yes	Yes	Yes	Yes	Yes	Yes	No
<b>Country-gender trend</b>	No	No	No	No	No	No	Yes
<b>Avg. dep. var.</b>	0.589	0.586	0.572	0.592	0.588	0.592	0.589
<b>Avg. (women)</b>	0.380	0.379	0.365	0.374	0.380	0.380	0.380
<b>Avg. (men)</b>	0.620	0.616	0.602	0.622	0.619	0.624	0.620
<b>Observations</b>	1.300e+08	1.320e+08	1.070e+08	1.160e+08	1.290e+08	1.260e+08	1.300e+08
<b>N. clusters</b>	640	642	518	636	604	220	640

Note: This table shows different robustness checks of the main specification on the relative effect of COVID-19 on women’s contributing activity. Each column (numbered at the top for easier reference) corresponds to the results of a different regression that captures a different robustness check. The sample is all the authors that created at least one commit during 2019 that were localized in a country and that were assigned a gender. Woman is a binary variable indicating if the author is classified as woman. The social distance index is a measure of the severity of the COVID-19 related social distance measures applied in each country each week. All columns include author and week fixed effect, as indicated by a “Yes” in the corresponding line and column. Column 1 is our baseline specification for comparison. Columns 2–7 apply different variations to the specification for robustness checks. Columns 1–6 include worldwide gender-specific trends. Columns 2 and 3 change the list of ccTLDs that are excluded from our geo-localization method: column 2 does not exclude any ccTLD nor any multi-localized company, while column 3 excludes additional ccTLDs (on top of the initial list) using the list from the “Country code top-level domains with commercial license” Wikipedia page (accessed December 2023). Column 4–6 restrict our sample by excluding some countries: column 4 exclude the USA; column 5 excludes countries that are insufficiently representative of their population. In particular, we sort countries by the ratio of authors appearing in our dataset of gendered and localized authors over the total population of the country, and we eliminate the top and the bottom 5%. Column 6 excludes countries with too few observations. In particular, we eliminate countries with less than 100 women (or men) authors. Column 7 shows the results when using country-specific gender-specific trends. Standard errors are clustered at the gender-country-post treatment level. Stars represent different levels of statistical significance: \*\*\* p<0.01; \*\* p<0.05; \* p<0.1.