NBER WORKING PAPER SERIES

TRANSGENDER EARNINGS GAPS IN THE UNITED STATES: EVIDENCE FROM ADMINISTRATIVE DATA

Christopher S. Carpenter
Lucas Goodman
Maxine J. Lee

Working Paper 32691
http://www.nber.org/papers/w32691

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2024

We are grateful to Olivia Compton, Marcus Dillender, Erwan Dujeancourt, Taryn Eames, Donn Feir, Gilbert Gonzales, Rob Haggar, Sam Mann, Mike Martell, Josh Martin, Laura Nettuno, Elisabeth Perlman, Jamie Ryan, Dario Sansone, Edson Severnini, Lawrence Stacey, Christina Sun, Lucas Tilley, Arthur Wickard, seminar and conference participants at Carnegie Mellon University (Heinz), ITAM, RAND, University of Illinois – Chicago, University of Maryland, Vanderbilt University, the CSQIEP Virtual Seminar on LGBTQ+ Economics, the 2024 Washington Area Labor Economics Symposium, and the 2023 Office of Tax Analysis Research Conference for helpful comments. This research would not have been possible without the help and foresight of Geoffrey Gee. All errors are our own. This research was conducted while Goodman was an employee at the U.S. Department of the Treasury. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily reflect the views or the official positions of the U.S. Department of the Treasury. Any taxpayer data used in this research was kept in a secured Treasury or IRS data repository, and all results have been reviewed to ensure that no confidential information is disclosed. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Transgender Earnings Gaps in the United States: Evidence from Administrative Data
Christopher S. Carpenter, Lucas Goodman, and Maxine J. Lee
NBER Working Paper No. 32691
July 2024
JEL No. J0,J10

ABSTRACT

We provide the first evidence on transgender earnings in the US using administrative data on over 55,000 individuals who changed their gender marker with the Social Security Administration and had gender-congruent first name changes on tax records. We validate and describe this sample which exhibits positive selection likely associated with the ability to legally affirm gender. To address selection we estimate transgender earnings gaps using timing variation within-person and variation across siblings and coworkers. All three approaches return evidence of robust transgender earnings penalties of 6-13 log points driven by extensive and intensive margin differences.

Christopher S. Carpenter
Department of Economics
Vanderbilt University
VU Station B, Box #351819
2301 Vanderbilt Place
Nashville, TN 37235
and NBER
christopher.s.carpenter@vanderbilt.edu

Lucas Goodman
Office of Tax Analysis
Department of Treasury
1500 Pennsylvania Avenue NW
Washington, DC 20220
Lucas.Goodman@treasury.gov

Maxine J. Lee
San Francisco State University
mclee@sfsu.edu

Transgender people – individuals whose current gender does not align with their sex assigned at birth – are a sizable share of the United States population.[1] Gallup data from 2023 indicate that 2.8 percent of Generation Z individuals (those born between 1997 and 2012) and 1.1 percent of Millennials (those born between 1981 and 1996) identify as transgender (Jones 2024). Flores et al. (2016) estimate that there are about 1.4 million transgender adults in the United States. In addition, transgender people are the disproportionate focus of social and policy discussions, including debates about whether transgender individuals should be able to access restrooms and play sports on teams consistent with their gender (as opposed to their sex assigned at birth), debates about whether transgender youths should be able to access gender affirming care, and debates about whether governments should require private businesses to treat transgender people the same as cisgender people. In one of the highest profile debates about transgender people in the labor market, a 2020 US Supreme Court decision in *Bostock vs. Clayton County* ruled that employment discrimination on the basis of transgender status is illegal under federal civil rights protections prohibiting sex discrimination. Yet despite this increasing focus on transgender people in the U.S., little is known about their economic outcomes.

There are many channels through which transgender status could be related to employment and earnings. First, transgender individuals may face discrimination in labor markets, and this may also extend to other settings such as housing (Abbate et al. 2024) and public accommodation, all of which may make it difficult for transgender people to secure stable employment and earnings (Center

[1] Individuals whose gender aligns with their sex assigned at birth are cisgender. Transgender women are individuals assigned male at birth who identify as women. Transgender men are individuals assigned female at birth who identify as men. There is a wide variance in the use of these labels; for example, 'transgender women' can be used by individuals who are assigned male at birth and identify as a woman but have not taken steps to change their gender expression. Not all transgender individuals desire to take medical, legal, and/or social steps to affirm their gender, and not all individuals whose gender and sex assigned at birth do not align identify as transgender.

for American Progress and Movement Advancement Project 2015, James et al. 2016). There is also evidence that transgender youth experience harassment, bullying, and discrimination in educational environments (Kosciw et al. 2014), leading to a lower likelihood of high school graduation and college attendance (Sansone 2019). Second, the specific health vulnerabilities and health needs of transgender individuals could affect their ability to work and/or their productivity at work. For example, transgender individuals have a higher likelihood of activity limitations, mental health conditions (such as clinical depression), and substance use disorders (Grant et al. 2011, James et al. 2016, McDowell et al. 2019, Branstrom and Pachankis 2020), which could be related to labor market opportunities. Challenges in accessing gender affirming care, especially in the current policy environment where states are increasingly banning access to such care for minors (Movement Advancement Project 2023), could also affect health capital as well as human capital acquisition and subsequent labor market opportunities. Finally, it is possible that earnings allow individuals to cultivate their gender identity (Akerlof and Kranton 2000) and access resources such as health insurance that covers gender affirming care.

Survey evidence suggests that transgender adults in the United States have lower employment and household incomes than similarly situated cisgender people (Badgett et al. 2021, Badgett et al. 2024 forthcoming, Carpenter et al. 2020, Carpenter et al. 2022, Shannon 2021, Stacey et al. 2022). The survey-based evidence is limited in important ways, however. First, the surveys have very small samples of self-identified transgender people. Second, existing surveys include questions on employment but generally lack information on labor market earnings, occupation, industry, or firm characteristics. Third, survey evidence on self-reported transgender status suffers from concerns about underreporting due to anti-transgender sentiment (Aksoy et al. 2024), cisgender people who purposefully misreport their gender in protest of questions inclusive of diverse gender identities, and other types of reporting biases, raising questions about the fidelity of existing estimates. Fourth, the surveys rely on self-reports of transgender status, which may not include those who do not identify as 'transgender' despite incongruence between their gender and sex assigned at birth.

In this paper we overcome the limitations of existing survey-based evidence on transgender status and economic outcomes to provide the first comprehensive evidence on earnings gaps for transgender people in the United States. To do so, we use confidential administrative data on individuals who changed their gender marker with the Social Security Administration and who also had gender-congruent first name changes on tax records with the Internal Revenue Service.[2] We identify over 55,000 likely transgender individuals, which is one to three orders of magnitude larger than any previous study in economics on transgender people.[3] We first document a variety of patterns in this sample to provide evidence that the gender marker changes reflect true gender affirmations, including the fact that the transgender sample has increased sharply over time and that transgender men are much younger than transgender women when they take administrative steps to affirm their gender. Both of these patterns align with demographic trends of self-identified transgender people in the United States from recent high-quality surveys (Kinzinger et al. 2023, Brown 2022, James et al. 2016). We also provide estimates of socioeconomic characteristics of transgender people in the United States using these administrative data, a first in the literature

[2] In section III.A, we describe in detail how we operationalize the idea of a 'gender-congruent' name change.

[3] For expositional ease, we refer interchangeably to 'likely transgender' and 'transgender' individuals when referring to our main analysis sample, although we recognize that some individuals who legally change their name and gender and take steps to medically and socially affirm their gender may not identify with the term 'transgender'. Conversely, we refer to 'cisgender' individuals when referring to those who have not changed their name or gender marker in government records, although this group includes some transgender individuals who have not taken steps to change their name and gender marker in government records.

that has heretofore exclusively relied on survey self-reports. We document strong positive selection in our transgender sample which is likely associated with the ability to legally and (in some cases) medically affirm one's gender.

To address selection, we take multiple complementary approaches to estimate transgender earnings gaps. First, we exploit panel variation, comparing earnings of transgender individuals before and after their name change to estimate within-person models of earnings in a panel evaluation framework. This approach eliminates time-invariant person-specific unobserved heterogeneity and returns a transgender earnings penalty of about 11 log points.[4] Second, we use the tax records to identify likely siblings of transgender people by linking individuals to the people who claimed them as dependents in their childhood (generally their parents) and finding other individuals in that same household who were also claimed as dependents. This allows us to exploit variation in transgender status within families across siblings for an alternative estimate of the transgender earnings gap. This sibling fixed-effects approach returns a transgender earnings penalty of 6 to 13 log points, depending on specification. Finally, we exploit variation across transgender and cisgender coworkers within the same firm and occupation. This approach also indicates a transgender earnings penalty of about 8 log points. Taken together, we conclude that the earnings gap associated with transgender status in the United States is robustly negative, in the range of 6 to 13 log points. If transgender people who are unable to take these legal and administrative steps to affirm their gender experience more severe discrimination or more vulnerable health outcomes compared to our sample, we expect the 6 to 13 log points earnings gap to represent a lower bound of the earnings penalty that transgender Americans experience on average.

[4] The timing of the individual's name change is the only individual transition we observe in the data. It is likely that social, medical, and legal changes preceded the administrative name change, and indeed in the event study models we show below there is evidence of an earnings decline before the individual's name change year.

I. Literature Review

Economists have long been interested in exploring wage differences related to gender, race/ethnicity, disability, parenthood status, sexual orientation, religion, national origin, indigenous status, geography, and a host of other demographic characteristics. Our research on transgender status is arguably most closely related to the economics literature on the gender wage gap (reviewed in Blau and Kahn 2017). The gender wage gap refers to the persistently lower wages for women compared to men. Despite changes in factors such as the gender difference in human capital accumulation, intrahousehold specialization patterns, and gender norms in society, there is a persistent gender wage gap that is not explained by observed differences in characteristics between women and men both in the US (Blau and Kahn 2017) and other developed countries (Kunze 2017). The gender gap is especially important to consider in a study of transgender population because labor market differentials between transgender and cisgender people are intrinsically linked to gender. For example, a transgender person assigned female at birth who has socially or medically affirmed that he is a man may earn higher income due to the gender premium for men or earn lower income due to a transgender penalty. Moreover, gender gaps in labor market outcomes are also affected by one's gender expression (Weichselbaumer 2001, Gorsuch 2019, Burn and Martell 2022), implying that transgender people's experience in the labor market may vary based on the specific medical or social steps taken to affirm their gender. Similarly, both the gender gap and transgender gap in potential earnings may change as individuals take steps to affirm their gender.

Evidence on labor market outcomes of transgender people is very limited in economics. Existing studies have taken two general forms: cross-sectional comparisons of economic outcomes for transgender people relative to demographically similar cisgender people (primarily using survey data in the US) and within-person comparisons of earnings of transgender people measured before and after their gender affirmation (using administrative data from the Netherlands).[5] Of the cross-sectional studies, Carpenter et al. (2022) use the US Census Bureau's nationally representative Household Pulse dataset – which asks individuals a question about sex at birth separately from a question about current gender – to document that transgender individuals have significantly worse economic outcomes (e.g., employment and personal income) than cisgender people. Stacey et al. (2022) find broadly similar patterns for self-identified transgender people compared to cisgender people in the Gallup Well Being Index data. Carpenter et al. (2021) and Mann (2021) also find similar patterns for self-identified transgender people compared to cisgender people in the 30+ states that asked a direct question about transgender status in the Centers for Disease Control and Prevention's Behavioral Risk Factor Surveillance System. None of these large, representative datasets in the US includes information on labor market earnings or firm characteristics.[6]

Elsewhere, one prior study in the Netherlands has used administrative records to identify transgender people and examine the effects of gender affirmation on earnings.[7] In the Netherlands (and other western European country

---

[5] Two recent resume experiment studies have tested whether resumes of fictitious job applicants that use 'they/them' pronouns generate differences in labor market outcomes as measured by callback rates for job interviews. These studies are relevant because some (but not all) individuals who use 'they/them' pronouns may identify as transgender and some (but not all) transgender individuals may use 'they/them' pronouns. Kline et al. (2022) examine jobs at Fortune 500 firms, while Eames (2024) studies this question in the context of job postings in Colorado, Utah, and Washington. While Kline et al. (2022) find a small penalty for resumes using 'they/them' pronouns (an effect on the order of about 1.5 percentage points), Eames (2024) finds a much larger penalty (about 5.5 percentage points) that is even larger in more conservative counties.

[6] Other studies use non-representative samples from surveys such as the 2015 United States Transgender Survey. Shannon (2021) and Campbell et al. (2023) find broadly similar patterns to the research relying on representative samples described above.

[7] One very small study in the US examined within-person changes in earnings of transgender people. Schilt and Wiswall (2008) studied 18 transgender women and 25 transgender men from a

contexts) where population and employment registers are linked to publicly funded medical care, transgender people can be identified through gender dysphoria diagnoses or a legal gender marker change. Geijtenbeek and Plug (2018) use administrative data from the Netherlands to identify 155 transgender men and 324 transgender women who have undergone gender affirmation surgery and sterilization; they find that transgender people are younger on average and have lower average earnings than cisgender people with the same sex at birth. They also find that earnings of likely transgender women fell after transition while earnings of likely transgender men increased.[8]

Our paper makes several important contributions relative to prior work. First, no prior study using survey or administrative data has examined transgender labor market earnings, occupation, industry, or firm characteristics in the United States. While the evidence on earnings from the Netherlands is useful, each country has markedly different attitudes and policies regarding gender minorities. For example, of the 175 countries ranked according to LGBTI acceptance in Flores (2021), the Netherlands ranks second in acceptance (behind only Iceland); in contrast, the United States ranked 23rd. Second, in addition to having strong confidence in the fidelity of our measurement of transgender status, our sample size of transgender people (approximately 55,000) is also much larger than used in prior work.[9] Relative to the only other panel data evaluation in the United

---

convenience sample of individuals in the US who attended transgender conferences or visited a transgender-related website. They found that earnings of transgender women fell after transition, while earnings of transgender men increased slightly.

[8] Another relevant study is Kolk et al. (2023) who examine individuals who received a diagnosis of gender incongruence (7,604 individuals) and/or who changed their gender on legal documents in Sweden (2,959 individuals) from 1973-2020. They document that transgender people have significantly worse socioeconomic outcomes than cisgender people with the same sex at birth, though they do not directly examine earnings nor do they exploit within-person timing variation.

[9] Our approach for identifying transgender people that relies on a combination of gender marker changes with the Social Security Administration and name changes on tax documents – described in detail in Appendix A – is closely related to Cerf Harris (2015) who linked SSA gender marker and name changes with 2010 Decennial Census data using the Census NUMIDENT file. That

States (Schilt and Wiswall 2008), our sample size is three orders of magnitude larger. Relative to recent European studies using administrative data, our sample size is two orders of magnitude larger. And relative to most survey-based studies in the US, our sample size is an order of magnitude larger. This allows us to make more precise inferences about transgender earnings gaps. Third, we are the first paper in the literature to link transgender people to their families and to their coworkers, which allows us to introduce new designs to estimate transgender earnings gaps, in addition to a within-person panel evaluation design.

II. INSTITUTIONAL CONTEXT

As we describe below in Appendix A, our primary analysis sample relies on individuals who have actively changed their gender marker with the Social Security Administration (SSA) and their first name in tax records. In this section we describe the process of changing one's first name and gender marker in government records and why transgender individuals change first name and gender marker with the SSA.[10] Often, individuals change both their first name and gender marker at once to reduce the number of times they need to amend government records.[11]

For a first name change, individuals file a petition with the court and are required to do some combination of paying a court fee, submitting to a background check at a local law enforcement agency, publishing a name change announcement in a local newspaper, and appearing before a judge. In the last decade, many states removed some of these requirements when the petitioner is

---

paper documented that likely transgender individuals (based on SSA gender marker changes) were significantly more likely to refrain from indicating a male or female sex and to mark both male and female sex on their 2010 Census form. Our research builds on this innovative administrative data effort to extend the analysis to labor market outcomes.

[10] We note that SSA gender markers are binary; there is no "X" option.

[11] In the subset of our data when the exact timing of gender marker and name changes can be observed, about 88 percent of likely transgender individuals appear to have changed their first name and gender marker at the same time. See Appendix B for further details.

changing the first name for gender affirmation. As of 2024, only nine states still require a name change announcement (Movement Advancement Project 2024).

For a gender marker change, the SSA required proof of gender affirmation surgery for a gender marker change until 2013 and a physician's letter certifying medical treatment for gender affirmation or a correct gender in other government-issued documents, such as the birth certificate or passport, from 2013 to 2022 (National Center for Transgender Equality 2013, 2024).[12] For those changing their gender marker on birth certificates, the process varied by state: some required individuals to obtain a court order, a physician's letter certifying gender identity, or proof of gender affirmation surgery, while a few states barred agencies from amending the gender marker on birth certificates. For changing the gender marker on passports, the Department of State required a physician's letter certifying medical treatment for gender affirmation from 2010 to 2021. Since 2021, documentation is no longer needed to change the gender marker on passports.

Once individuals obtain both the court order for first name change and one of the forms that certify the gender identity, they submit the application for a replacement social security card to reflect these changes in the SSA records. This is a critical step in changing the first name on driver's licenses or identification cards in some states that verify the changed name with the SSA records. Also, employers often verify an individual's name with the SSA records for tax purposes. Individuals who go through the name change process to update their name with the SSA for these reasons would also update their gender marker using the same form. Those receiving Medicare or Medicaid benefits may be further

[12] Since 2022, SSA has allowed individuals to change their gender marker by simply requesting a replacement social security card with proof of identity, which does not need to have a gender marker that matches the requested gender marker. This recent change is mostly outside our analysis window, however.

incentivized to amend their gender marker because they may be denied coverage for some medical treatments based on the gender listed in the SSA records.

There is little evidence on the share of people who identify as transgender who take legal steps to affirm their gender. None of the representative surveys described above asks such questions, for example. Studies of non-representative samples such as the 2015 and 2022 United States Transgender Surveys generally suggest that about 40-50 percent of transgender adults have at least one identity document with their affirmed gender (James et al. 2024, Herman and O'Neill 2020, Restar et al. 2020). As the process to change the first name and gender marker is often expensive and/or complex, we expect our sample to consist of a positively selected subset of the transgender population.

III. Describing Transgender Individuals in US Administrative Data

In this section we describe the characteristics of transgender individuals in our administrative data. Appendix A contains a detailed description of the two main datasets we use to identify transgender individuals, as well as how we reach our final sample and confirm that the sample we identify is very likely to include transgender individuals who have affirmed their gender. To summarize, we observe the universe of individuals whose Social Security Administration (SSA) gender marker was changed from male to female or from female to male since 2011. We link these individuals to tax records from the Internal Revenue Service and identify likely transgender individuals as those who not only have a SSA gender marker change but also have a gender-congruent first name change in the tax records, in the spirit of Fisher, Gee, and Looney (2018). For each individual in our sample we construct a panel from 2000 through 2022. Our primary outcome of interest is earnings, which is the sum of wage income (which we retrieve from

Form W-2) and self-employment income (which we retrieve from Schedule SE of Form 1040).[13]

Appendices A and B describe in detail the requirements we impose for first name changes to meet our sample inclusion criteria, as well as the likely implications of each of our choices for sample representativeness. For example, we show that the vast majority (almost two thirds) of individuals with an SSA gender marker change exhibit no first name change in tax records. While some of these individuals may in fact be transgender and identify as such, we exclude these individuals from our analysis sample and show in the Appendix B that their characteristics are significantly different from those with gender-congruent name changes (i.e., individuals we identify as likely transgender). Although our sample restrictions result in a non-representative sample of transgender individuals, we describe multiple complementary approaches designed to isolate plausible estimates of transgender earnings gaps where these selection concerns are substantially reduced or eliminated.

Figure 1a graphs trends over time in the number of individuals we identify as transgender, where an individual is assigned to a calendar year based on the timing of their name change in the IRS Form 1040 available from 2008 to 2022. We refer to this calendar year as the "name-change year" rather than "transition year", since affirming one's gender identity and/or gender expression is a process that may begin several years prior to (and continue for several years after) a name change.[14] We present graphs separately for transgender women (closed circles)

[13] Very similar rates of cisgender and transgender individuals – around 3-4 percent – have self-employment income as their primary source of income.

[14] In particular, the name-change year is (in most cases) equal to the year of the Form 1040 in which the post-transition name is first used. However, if an individual does not appear on Form 1040 (as a primary filer, secondary filer, or dependent) from time $t-k$ to time t for $k>1$, we impute a name-change year randomly (uniformly) among $t-k+1, t-k+2, \ldots, t$. Additionally, in Appendix B, we show how the name-change year (measured as described in this section) relates to the identified dates of record changes for a subset of the sample where the latter can be observed.

and transgender men (open circles). Figure 1a documents that we identify increasing numbers of transgender individuals over time, consistent with survey evidence that the likelihood of self-identifying as transgender has increased significantly over time (e.g., Jones 2024). Figure 1a additionally shows a reduction in name changes in 2020. This may reflect extended COVID-19-related closures of SSA offices, which likely made it costlier for individuals to update their name or gender marker with the SSA.

Figure 1b presents the age profile of our transgender sample, again separately for transgender women (closed circles) and transgender men (open circles), measured in the name-change year. Two patterns are clear from Figure 1b. First, the transgender sample is disproportionately young. Most individuals we identify as transgender change their names in their late teens and early 20s. Second, the sample of transgender men has name-change ages that are younger than the sample of transgender women. These patterns are consistent with recent credible survey evidence on the demographics of the transgender population (Carpenter et al. 2022).[15]

Figures 2a and 2b show the geographic distribution of the transgender men and transgender women in our sample, respectively, measured in 2022.[16] Each figure shows a map where more darkly shaded states represent higher shares of each state population that is composed of transgender people. Figures 2a and 2b show that transgender individuals are systematically more prevalent in the more

[15] The young age of the sample and the increasing time series pattern are also consistent with the samples of transgender individuals identified in health claims data (McDowell et al. 2019, Baker and Restar 2022).

[16] The state averages are expressed as a share of the population with the same sex assigned at birth; that is, the transgender women's share is the ratio of transgender women to the sum of cisgender men and transgender women. We define each individual's residential location by the zip code reported on their Form 1040; for non-filers, we use the most common zip code reported in a given year on information returns (such as Form W-2) sent to that individual. If an individual is neither a filer, dependent, nor a recipient of an information return in a given year, we impute geography forward from the most recent non-missing year.

liberal coastal parts of the United States as well as in Colorado and Minnesota, and substantially less prevalent in the South. The geographic distribution of transgender women and transgender men is very similar; the state shares of transgender women and transgender men have a correlation coefficient of 0.95. Our implied state-level population shares (aggregated between transgender women and transgender men) are also strongly positively correlated with the self-identified transgender shares across states from the Census Bureau's Household Pulse dataset (with a correlation coefficient of 0.76).[17]

Table 1 presents means for a range of outcomes that can be meaningfully identified from the IRS data. The format of Table 1 is as follows: we present means for transgender men in Column 1, cisgender men in Column 2, transgender women in Column 3, and cisgender women in Column 4.[18] To account for the systematically younger age distribution of our transgender sample shown in Figure 2, the means for cisgender men (women) have been reweighted to match the year and cohort distribution of the transgender men (women). Table 1 indicates that transgender individuals live in counties that vote slightly more Democratic and live in zip codes where the share of the population that is Black is

[17] Notably, we find substantially more variation in state-level shares in our data than in Household Pulse; the coefficient of variation of state population shares is 0.77 in our data and 0.14 in Household Pulse. Put differently, both our data and Household Pulse identify Oregon as the state with the largest transgender share, but 6.0% of our sample lives in Oregon, while the Household Pulse estimates imply that 2.1% of transgender people in the US live in Oregon. This difference may be caused by heterogeneity in state-level policies that factor into the administrative burden of obtaining a court order for changing first name or gender marker and changing gender markers on a birth certificate, which could be used to change one's SSA gender marker for much of our sample period.

[18] Specifically, we form a panel out of a "control" sample designed to be approximately representative of the U.S. population. The cisgender samples are those individuals who are in the control sample and not also in the transgender sample. We describe the construction of this control sample in Appendix C. We note that some portion of the "cisgender sample" may not be cisgender. All individuals appearing on a tax return or receiving an information return who are not in the transgender sample – including transgender individuals who have not changed SSA gender marker – are in the universe of this sample.

slightly smaller. Regarding education,[19] Table 1 indicates that our transgender sample is very positively selected: 87.5 percent of transgender men and 87.6 percent of transgender women have any college education, compared with only 64.3 percent of cisgender men and 75.8 percent of cisgender women.[20] Regarding marital and family outcomes, transgender people are much less likely to be married than cisgender people, but they are much more likely than cisgender people to be married to someone we identify as transgender.[21] Table 1 also shows that transgender people are less than half as likely as cisgender individuals to have dependents.[22]

We also describe family background characteristics for further evidence on the presence and degree of positive selection. This is possible due to the parent-child links in the IRS tax records where we observe the individuals who claimed transgender individuals as dependents when those transgender people were children (this process is described in detail in Appendix D). The first row of the bottom panel of Table 1 presents the likelihood that we are able to link an individual to their parent, conditional on the reference individual being born in

---

[19] We measure years of education as the number of distinct calendar years in which an individual receives Form 1098-T, which reports tuition payments and scholarships. We observe Form 1098-T receipt beginning in 1999. We do not observe whether an individual successfully attained a degree. The education statistics are conditioned on having year of birth between 1981 and 1999 – i.e., attaining age 18 no earlier than 1999 and attaining age 23 by 2022 (the end of our sample).

[20] Notably, this positive selection is not found in large surveys. In the Census Bureau's Household Pulse dataset, for example, transgender men and women have lower education and income than cisgender individuals with the same sex at birth (Carpenter et al. 2022).

[21] Note that cisgender men and cisgender women are not directly comparable because they are reweighted to samples with very different age distributions. That is, we reweight the cisgender men to have the same age distribution as transgender men but we reweight the cisgender women to have the same age distribution as transgender women. As transgender men are significantly younger than transgender women on average (as shown in Figure 1b), this partly contributes to the lower college education and marriage rates of cisgender men than cisgender women in Table 1.

[22] Carpenter et al. (2024, forthcoming) similarly find that transgender and gender diverse adults in the Household Pulse are less likely to live with any person under 18 at home compared to cisgender people. The magnitude of the difference is much smaller in Carpenter et al. (2024, forthcoming) potentially due to the Household Pulse asking for the number of people under 18 living in the household rather than each person's dependent status.

1977 or later (to match the years that we have good parent-child linked data coverage). Interestingly, and consistent with the positive selection on education, our transgender samples are slightly more likely to have a valid parental link than the cisgender samples. Evidence of positive selection is even more clear in the comparison of average parental incomes in 2022 dollars: we see that transgender men (women) were in households with parental incomes approximately $116,000 ($136,000) in 2022 dollars, while cisgender people were raised in households with parents who earned about $85,000. Table 1 also shows that the likelihood the individual grew up in a single parent household is much lower for the transgender sample than for the cisgender sample and that transgender individuals grew up in places that had higher White shares in the individual’s zip code. Overall, these patterns indicate that the individuals we identify as transgender grew up in households with significantly greater material advantage than cisgender individuals. These facts motivate our approaches to estimate transgender earnings gaps that can plausibly address this positive selection.

IV. Estimates of Transgender Earnings Gaps

In this section we present evidence on transgender earnings gaps from three complementary approaches designed to isolate plausibly causal estimates of the transgender earnings gap: a within-person panel approach and two cross-sectional approaches that compare transgender and cisgender siblings and coworkers, respectively.

A. Panel evidence

Our first approach exploits variation in the timing of a transgender individual’s name change, which Appendix E documents is associated with sharp changes in medical expense deductions, cross-state moves, and marital status changes. In the years surrounding the name change, transgender individuals may come out to more of their friends and colleagues, inform their employers of their gender identity (e.g., through gender and name changes in the administrative records), take time off from work for medical care, and/or change their presentation to match their affirmed gender. Within-person changes in earnings identify earnings gaps experienced by transgender individuals under a set of identification assumptions: (a) a standard parallel trends assumption that other factors affecting earnings are not changing systematically in the vicinity of the name-change year and (b) individuals' earnings sufficiently prior to their name-change year are as if they were cisgender, with a gender identity equal to their sex assigned at birth. While we support assumption (a) by analyzing differential trends prior to name-change, we expect assumption (b) not to hold strictly, implying that the true gap is likely larger in magnitude than what is uncovered by this panel analysis.

We examine within-person changes in earnings for transgender individuals, separately by gender. We use cisgender individuals of the same sex assigned at birth (in line with assumption (b) above) as controls. We restrict to those transgender individuals whose age at name-change is at least 28, so that pre-transition earnings are meaningful, and whose name-change year is 2018 or earlier, so that we observe earnings for at least four years after the name-change year. We note that these are significant restrictions: the age restriction drops 49% of transgender women and 74% of transgender men, while the name-change year restriction drops an additional 25% of transgender women and 13% of transgender men.

We use a stacked event study approach. For each cohort of transgender individuals (defined by their name-change year τ), we form a panel from event times -6 to +4. We append a control panel for τ, consisting of all cisgender individuals in the control sample of the same sex assigned at birth, using observations from $t = [\tau - 6, \tau + 4]$; we reweight the samples so that they have the same age distribution. We repeat this data construction for all cohorts τ.[23] With this data construction, we run an event study regression (separately by gender) with earnings as the relevant outcome, controlling for person fixed effects (interacted with cohort) and event time dummies (again interacted with cohort) and omitting the event study coefficient for -4. We choose to omit event time -4, rather than -1, reflecting the fact that gender affirmation is a multi-year process that may begin several years prior to the name-change year. In the language of assumption (b), we are assuming that at event time -4, a transgender individual has earnings as if they were a cisgender individual with gender identity corresponding to their sex assigned at birth.

The conditional expectation function for earnings is plausibly exponential i.e., covariates such as time fixed effects likely affect earnings in a proportional, not additive, manner. Yet, individual observations have zero earnings with non-trivial frequency. These two facts motivate us to follow recent literature and estimate a Poisson regression rather than using OLS with log earnings as the dependent variable (Chen and Roth 2024).[24] In particular, the Poisson regression estimates the following model for person i at time t whose cohort (as described above) is τ, whose sex assigned at birth is $SAAB_i$ and whose gender identity is G_i:[25]

$$\log\big(E(y_{it\tau}|G_i, SAAB_i)\big) = \lambda_{i\tau} + \mu_{t\tau} + \sum_{e=-6, e\neq -4}^{4} \beta_e \times 1(t = \tau + e) \times 1(G_i \neq SAAB_i) \tag{1}$$

[23] This data construction means that the same control individual may appear in multiple panels. We cluster our standard errors by unique individual, which accounts for any mechanical correlation of residuals created by such duplication.
[24] We use the Stata command ppmlhdfe (Correia, Guimaraes, Zylkin (2020)) to estimate these regressions.
[25] G_i in this regression is defined to be time-invariant (female for transgender women and cisgender women, and male for transgender men and cisgender men).

The resulting coefficient estimates reflect the difference in average earnings (in log points) for transgender individuals from event time -4 to other event times, relative to the difference for similar-age cisgender individuals (of the same sex assigned at birth) from event time -4 to other event times.[26]

In Figure 3, the red series with the dashed line reports the results of this regression for transgender women, using cisgender men as controls. We find that earnings are relatively flat from event times -6 through -3, suggesting that earnings during this period may, once adjusted for event time fixed effects, serve as a plausible counterfactual for post-name-change wages. Beginning at event time -2, earnings fall for transgender women relative to cisgender men. By event time +4, earnings stabilize, with a total gap (relative to event time -4) of -17.9 log points, with a standard error of 2.6 log points.[27] This total effect of -17.9 log points reflects an estimate of the effect of transitioning for transgender women. The black solid line repeats the same event study regression for transgender men, using cisgender women as controls. We find relatively flat trends prior to event time -4, followed by a modest decline of about 3.7 log points.

The 17.9 and 3.7 log point effects reflect estimates of the effect of transitioning for transgender women and men, respectively. However, there are multiple ways of interpreting these two effects. Under the assumptions underlying the panel approach, the earnings of a transgender woman at event time +4 differs from herself at event time -4 both in that she is perceived as transgender and that she is perceived as a woman. Thus, the -17.9 log point effect may reflect both a penalty for being transgender and a penalty for being a woman. Likewise, the -3.7

[26] The coefficient estimates from this regression could approximately be estimated by collapsing mean earnings to the transgender-by-event-time level, and then running the (exactly-identified) event study OLS regression on the natural log of mean earnings. In Appendix Figure I1, we plot log mean earnings separately by transgender men, transgender women, cisgender men, and cisgender women.

[27] We report all event study coefficients discussed in this section, along with their standard errors, in Appendix Table I2.

log point effect for transgender men may reflect both a penalty for being transgender and a bonus for being a man.

Formally, let $\alpha^{G,SAAB}$ denote the change in log mean earnings from event time -4 to event time +4 for a given group defined by gender and sex assigned at birth. One can write $\alpha^{G,SAAB}$ as:

$$\alpha^{G,SAAB} = \beta^{F,F} + \alpha_{male} \times (G = M) + \alpha_{MAAB} \times (SAAB = M) + \alpha_{trans} \times (G \neq SAAB) \tag{2}$$

The α terms on the right-hand-side can be interpreted as the direct effect of male gender, male sex-assigned-at-birth, and transgender status, respectively; these four terms (including the constant) are pinned down exactly by the four values of $\alpha^{G,SAAB}$ estimable from the data. In the context of this model, our panel evidence estimates $\alpha_{trans} - \alpha_{male} = -0.179$ and $\alpha_{trans} + \alpha_{male} = -0.037$. These estimates imply an estimate for α_{trans} equal to -0.108 (i.e., the average of -0.179 and -0.037), with a standard error of 0.016.

We interpret this -10.8 log point gap as an estimate of the residual earnings penalty experienced by transgender individuals even after accounting for time invariant person-specific factors that may cause productivity to vary. We provide two cautions for the interpretation of this estimate. First, while the event study pictures suggest that earnings are fairly stable outside of event times -4 to +3, our estimate would understate the transgender penalty if transgender individuals were already experiencing differential treatment in the labor market at event time -4, or if the penalty continued to grow past event time +4. Second, our estimate includes any productivity-reducing effects caused by medical procedures or medication involved in gender affirmation; thus, the estimated penalty may reflect factors other than discrimination.

We also highlight that there are other possible interpretations of the two event study effects. For instance, it could be the case that the transgender gap itself is heterogeneous by gender. Relatedly, it is possible that the gender wage gap differs between transgender men and women – i.e., that the male bonus for transgender men does not equal the female penalty for transgender women. Both of these interpretations are consistent with possibilities that transgender men are more able to "pass" (that is, be perceived by employers and/or customers as cisgender men) or that women perceived to be transgender face more scrutiny and hostility than men perceived to be transgender. This more flexible model, which is not identified empirically, would be written as:

$$\begin{aligned}\alpha^{G,SAAB} = \beta^{F,F} &+ \alpha_{male,M} \times (G = M, SAAB = M) \\ &+ \alpha_{male,F} \times (G = M, SAAB = F) \\ &+ \alpha_{MAAB} \times (SAAB = M) \\ &+ \alpha_{trans,M} \times (SAAB = F, G = M) \\ &+ \alpha_{trans,F} \times (SAAB = M, G = F)\end{aligned} \tag{3}$$

The event study estimates would map to $-0.179 = -\alpha_{male,M} + \alpha_{trans,F}$ and $-0.037 = \alpha_{male,F} + \alpha_{trans,M}$. The transgender gap from Equation (2) of -0.108 represents the average of $\alpha_{trans,M}$ and $\alpha_{trans,F}$, under the assumption that $\alpha_{male,F} = \alpha_{male,M}$.

But the event study estimates are also consistent with many other (infinite) combinations of parameters in this more flexible model. For example, suppose the transgender effect itself is fairly small, but the labor market operated under a version of a "one-drop" rule, whereby all individuals whose sex assigned at birth *or* gender identity is female are treated analogously. This would require $\alpha_{male,F} = 0$; if we additionally imposed $\alpha_{trans,F} = \alpha_{trans,M}$, then we would recover $\alpha_{trans,F} = \alpha_{trans,M} = -0.037$ $\alpha_{male,M} = 0.142$. Under this alternate interpretation, the -10.8 log point transgender penalty primarily from Equation (2) reflects the fact that all transgender people (as opposed to only some cisgender people) are treated as women. Nevertheless, regardless of interpretation, the transgender penalty estimated in Figure 3 is a new and robust finding that transgender people have lower earnings, on average, after taking steps to affirm their gender.

Both the intensive (such as workplace treatment) and extensive (such as discrimination in hiring) margin differences can cause the transgender earnings penalty. Figure 4 shows that the extensive margin likely plays a significant role in driving the 11 log point transgender earnings gap. Figure 4 plots a stacked event study regression with a dummy for having any earnings as the dependent variable (i.e., using the right-hand-side in Equation (1)).[28] Between event times -4 and +4, transgender women's extensive margin participation falls by 9 percentage points relative to same-aged cisgender men; transgender men also experience a 6 percentage point drop in extensive margin participation relative to cisgender women. In Appendix F, we find that the extensive margin effect explains approximately half of the total penalty for transgender women, and over 100% of the (much smaller) penalty for transgender men.[29]

B. Cross-sectional Evidence Using Sibling Comparisons

To complement the panel-based evidence, we additionally make cross-sectional comparisons between transgender individuals and their cisgender siblings and coworkers. These approaches allow us to identify the transgender gap for younger cohorts, many of whom were too young to meet the necessary sample restrictions imposed in the panel approach. To implement the siblings approach, we define two people to be siblings if they (a) appear as dependents jointly on the same tax return at least three times and (b) are within five years of age. Because

[28] We estimate this regression using the user-written Stata command reghdfe (Correia 2016).

[29] We find no evidence the increase in non-participation is related to unemployment insurance, retirement income, or other capital income. We also see no increase in the likelihood of being a student, being claimed as a dependent, or being married while not working. We do see a quantitatively modest role for Social Security Disability Income (SSDI): the share of transgender women not working and receiving SSDI increases by 1.3 percentage points between event time -4 and 4 (relative to the age-matched cisgender men serving as controls). See Appendix F for further details.

we observe dependent linkages from 1994 onward and because most people are claimed as dependents until at least age 17, we successfully match most siblings if they are both from the 1977 cohort or later. We organize our data into two observations per sibling pair – one for the transgender individual and one for the cisgender sibling.[30] As our sample is fairly young, we use observations only from calendar year 2022, the latest available year, in order to capture earnings at the oldest possible age.

We make several other sample restrictions. First, we restrict to the 1981-1999 cohorts. We drop cohorts 1980 and older because we use years of post-secondary education as a control in some specifications; data limitations imply that we might miss the first years of education for cohorts prior to 1981. Likewise, we drop cohorts 2000 and younger, as they are in their prime college-going years in 2022. Finally, we require that both the transgender individual and the cisgender sibling meet all sample restrictions to be included in the regression. We include fixed effects for each sibling pair, and we cluster our standard errors at the household level (that is, at the level of the transgender member of the sibling pair). We again use Poisson regression with earnings as the dependent variable to account for the non-trivial presence of zeros.

By construction, two siblings will tend to have experienced similar levels of childhood resources and, additionally, tend to have similar genetic endowments. Nevertheless, as we show in Appendix Table I2, there remain substantial differences in observable characteristics across the siblings in adult outcomes such as marital status and childrearing; these remaining differences could easily drive substantial, systematic differences in earnings between transgender individuals and their siblings. Thus, we include varying sets of covariates in these regressions in order to recover the residual transgender gap.

[30] If a given transgender individual has k siblings in our sample, they appear k times in the data, and all observations in a pair including that individual are assigned a weight $\frac{1}{k}$.

We consider two sets of covariates. Broadly following the gender wage gap literature, we first consider a limited set of covariates related to human capital. Specifically, this set of covariates includes education controls (a linear term in the number of calendar years of college education and a dummy for currently being a college student, each constructed using Form 1098-T) and a quadratic in potential experience (age minus education minus 18). We also consider specifications where we add a second set of covariates related to family status (a dummy for being married, a dummy for having any dependent, and an additional dummy for having one's youngest dependent age 5 or younger, each constructed using Form 1040),[31] industry (fixed effects for two-digit NAICS codes, including missing), occupation (fixed effects for two-digit occupation, including missing), and geography (a linear control for the log population density of the zip code and a dummy for being in a central county in one of the top 30 MSAs ranked by population).[32] These latter covariates may serve as a proxy for other labor-market-relevant attributes (such as preferences or skills), but they may also be caused by labor market discrimination. We present results with both sets of controls for completeness.

Column 1 of Table 2 presents the baseline results without any controls beyond fixed effects for the sibling pair. Row 1 restricts to pairs of siblings consisting of a transgender woman and a cisgender man; we find that transgender

---

[31] For non-filers, we impute dependents and marital status forward from the most recent year in which the individual filed a tax return.

[32] We discuss in Appendix G how we measure industry and occupation, and Appendix Tables G1 and G2 show the industry and occupation distribution of transgender individuals, respectively. In many cases, such as in occupations and industries related to health care, we find that transgender individuals appear to have shares between those of cisgender men and women – that is, transgender women are overrepresented relative to cisgender women and transgender men are underrepresented relative to cisgender men, or vice versa. Arguably most strikingly, we also find that 17.2% of transgender women are in computer and mathematical occupations, compared to 3.9% of transgender men, 1.8% of cisgender women, and 4.3% of cisgender men. Within this occupation code, approximately half of the mass for transgender women is contained in the "Software Developers, Applications" occupation, consistent with anecdotal evidence discussed in Kychenthal (2022).

women earn 17.3 log points less than their cisgender brothers. Row 2 restricts to pairs of siblings consisting of a transgender man and a cisgender woman; we find that transgender men earn 1.7 log points more than their cisgender sisters, which is insignificantly different from zero. Our preferred interpretation, following Equation (2), deems the simple average of row 1 and row 2 as the transgender gap. In fact, we can improve the precision of this estimate by bringing in the full set of sibling pairs (i.e., including sibling pairs that do not match on sex assigned at birth) and including dummies for male gender identity and male sex assigned at birth to the regression. These results are presented in row 3: we estimate a transgender earnings gap of 7.9 log points. In the remaining columns, we add successively more detailed controls; we find that, if anything, the transgender earnings gap becomes more negative when controls for education, experience, family characteristics, industry and occupation dummies, and geography are included in the regression.[33]

C. Cross-sectional Evidence Using Coworker Comparisons

Finally, we present estimates of transgender earnings gaps using variation in transgender status across coworkers in the same occupation in a manner very similar in spirit to the siblings fixed-effects estimates. In particular, we match transgender individuals to a cisgender colleague who works at the same firm (based on the Employer Identification Number reported on their highest-wage Form W-2 in 2022) and reports the same occupation, at the three-digit level. A key advantage of this approach is the ability to eliminate unobserved firm-specific heterogeneity that may contribute to transgender earnings gaps such as firm climate toward gender minority people. This approach also allows us to compare

[33] In Appendix H, we report additional information on the role of each control. Specifically, we split the mentioned set of control variables into additional columns, in order to isolate their role more finely. Additionally, we consider a specification that includes controls for certain labor market variables, such as tenure at a given employer, in order to shed more light on some mechanisms driving this gap. We find that marriage and family controls tend to reduce the transgender gap, while geography and industry/occupation controls tend to increase it.

outcomes for transgender and cisgender workers of the same reported occupation within the same firm, and it sidesteps concerns about possible differential treatment of transgender and cisgender siblings by parents or other family members. A key sample restriction for this analysis is that we require individuals to be working with at least one year of job tenure, which means that the differences we estimate in the coworker analysis are inherently intensive margin differences. As in the siblings analysis, we restrict to individuals in the 1981-1999 birth cohorts. In sum, this dataset comprises of 186,030 pairs of 13,230 unique transgender individuals; we weight the data such that each transgender individual has the same aggregate weight.

Although this approach eliminates firm-specific differences, Appendix Table I3 shows that there remain some differences in observable characteristics between transgender and cisgender coworkers, including the fact that transgender individuals obtain more years of education than cisgender coworkers. Thus, as in the siblings analysis we include covariates in these regressions. We include fixed effects for each coworker pair and cluster our standard errors at the level of the transgender member of the coworker pair. Although mechanically this sample has positive earnings in 2022, we use Poisson regression with earnings as the dependent variable to maintain consistency with other specifications in the paper.

Table 3 presents results for this specification, following the format of Table 2. The first column does not include any controls beyond fixed effects for each coworker pair, and the remaining columns show sensitivity to inclusion of controls for education, family characteristics, and geography (note that occupation and industry are inherently fixed since we are comparing coworkers with same occupation working for the same employer). As with the siblings analysis, we find that transgender women earn less than their cisgender men coworkers – in this case, by about 11 log points. We also find that transgender men earn about 3.5 log points less than their cisgender women coworkers. When we combine all of our coworker pairs regardless of sex at birth we find a transgender earnings penalty of approximately 8 to 9 log points; we additionally find a 3 log point bonus to having male gender and an additional 6 log point bonus to being assigned male at birth. When we include control variables for education, potential experience, family status, and geography, the transgender women penalty decreases slightly and the transgender men penalty increases slightly; the overall transgender gap remains constant.

D: Summary of Estimates of Transgender Earnings Gaps

In sum, across three complementary approaches for addressing positive selection associated with our sample of transgender individuals, we consistently estimate a 6 to 13 log point residual penalty for transgender individuals, relative to their similarly situated cisgender counterparts. Each approach relies on a different sample construction and set of assumptions and limitations. Throughout, we emphasize that we are not interpreting the residual differences between transgender and cisgender individuals as solely being attributable to direct discrimination, though the estimates certainly include discrimination. For example, the effects of gender-affirming medical treatment may contribute to the residual earnings penalty, despite not being (directly) caused by discrimination. Nevertheless, we find it quite compelling that the two cross-sectional approaches yield qualitatively similar results to the within-person changes studied in the panel approach, despite the non-overlapping set of identification restrictions required.

VI. Conclusion

We use confidential administrative data to identify over 55,000 likely transgender individuals in the United States. Relative to survey-based samples, our sample is measured with very high fidelity: we identify individuals who actively changed their gender marker with the Social Security Administration and who changed their first name in a manner consistent with the gendered nature of the gender marker change (i.e., from relatively male to relatively female or vice versa) from 2009 to 2022. Using three complementary approaches – a within-person panel evaluation design, a within-family sibling fixed-effects design, and a within-occupation-and-employer coworker fixed-effects design – we find clear evidence that transgender status is penalized in the labor market, an effect on the order of 12 log points of annual earnings in the panel and 6 to 13 log points in the cross-section.

Our results are subject to some limitations, many related to the data. First, the sample we identify is not a random sample of transgender people, despite being a high fidelity near universe of individuals who actively changed their gender marker with the Social Security Administration and who had gender-consistent first name changes from 2009-2022. Not all transgender people take the steps to affirm their gender in government documents, and little is known about how doing so is related to demographic characteristics of transgender people. Second, our approach misses individuals who change their gender marker but do not change their name, individuals who rarely file Form 1040 even if they change their SSA gender marker, individuals who are unable to afford the administrative, legal, and/or medical steps required to change their name and gender marker in government records, and individuals who migrated to the US after 2011. Third, we do not observe medical records that would allow us to examine diagnoses or use of gender affirming care around the timing of the individual's name change. More generally, we do not know how the process of social transition maps onto the timing of one's name change, though the time path of relative earnings provides some insight.

Despite these limitations, our paper provides the first high quality evidence on the earnings differences associated with transgender status in the United States, a country with substantial heterogeneity in policies and attitudes toward transgender people. Overall our results are consistent with the idea that transgender people in the US may benefit from stronger labor market protections, as we note that the sample we identify is positively selected on own education and family background, especially with respect to parental income. This positive selection – which should not bias our design-based estimates of the earnings differences associated with transgender status – suggests that transgender people without such access to resources are likely experiencing worse economic outcomes than those we document here. More research is needed to understand how contextual factors shape the economic opportunities and outcomes of this growing and increasingly relevant population.

REFERENCES

Abbate, Nicolas, Ines Berniell, Joaquin Coleff, Luis Laguinge, Margarita Machelett, Mariana Marchionni, Julian Pedrazzi, and Maria Florencia Pinto (2024). "Discrimination against gay and transgender people in Latin America: A correspondence study in the rental housing market," *Labour Economics*, 87: article 102486.

Agresti, Alan and Brian Caffo (2000). "Simple and Effective Confidence Intervals for Proportions and Differences of Proportions Result from Adding Two Successes and Two Failures," *The American Statistician*, 54(4): 280-288.

Akerlof, George A. and Rachel E. Kranton (2000). "Economics and Identity," *Quarterly Journal of Economics*, 115(3): 715-753.

Aksoy, Billur, Christopher S. Carpenter, and Dario Sansone (2024, forthcoming). "Understanding Labor Market Discrimination Against Transgender People: Evidence from a Double List Experiment and a Survey," *Management Science*, forthcoming.

Badgett, M. V. Lee, Christopher S. Carpenter, and Dario Sansone. (2021). "LGBTQ Economics." *Journal of Economic Perspectives*, 35(2): 141-170.

Badgett, M. V. Lee, Christopher S. Carpenter, Maxine J. Lee, and Dario Sansone (2024, forthcoming). "A Review of the Economics Literature on Sexual Orientation and Gender Identity," *Journal of Economic Literature*, forthcoming.

Baker, Kellan and Arjee Restar (2022). "Utilization and Costs of Gender-Affirming Care in a Commercially Insured Transgender Population," *The Journal of Law, Medicine & Ethics*, 50(3): 456-470.

Blau, Francine D. and Lawrence M. Kahn (2017). "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, 55(3): 789-865.

Branstrom, Richard and John E. Pachankis (2020). "Reduction in Mental Health Treatment Utilization Among Transgender Individuals After Gender-Affirming Surgeries: A Total Population Study." *American Journal of Psychiatry*, 177(8): 727-734.

Brown, Anna (2022). "About 5% of young adults in the U.S. say their gender is different from their sex assigned at birth." *Pew Research Center*. Available at: https://www.pewresearch.org/short-reads/2022/06/07/about-5-of-young-adults-in-the-u-s-say-their-gender-is-different-from-their-sex-assigned-at-birth/. Last accessed December 5, 2023.

Burn, Ian and Michael E. Martell (2022). "Gender typicality and sexual minority labour market differentials." *British Journal of Industrial Relations*, 60(4): 784-814.

Campbell, Travis, Samuel Mann, Duc Nien Nguyen, and Yana Rodgers (2023, forthcoming). "Hormone Therapy, Suicidal Risk, and Transgender Youth in the United States," *AEA Papers and Proceedings*, forthcoming.

Carpenter, Christopher S., Samuel T. Eppink, and Gilbert Gonzales (2020). "Transgender Status, Gender Identity, and Socioeconomic Outcomes in the United States," *Industrial & Labor Relations Review*, 73(3): 573-599.

Carpenter, Christopher S., Maxine J. Lee, and Laura Nettuno (2022). "Economic outcomes for transgender people and other gender minorities in the United States: First estimates from a nationally representative sample," *Southern Economic Journal*, 89(2): 280-304.

Carpenter, Christopher S., Maxine J. Lee, and Laura Nettuno (2024, forthcoming). "Gender Minority Status and Family Inequality in the United States," *Russell Sage Foundation Journal of the Social Sciences*, forthcoming.

Cerf Harris, Benjamin (2015). "Likely Transgender Individuals in U.S. Federal Administrative Records and the 2010 Census," working paper.

Chen, Jiafeng and Jonathan Roth (2024). "Logs with zeros? Some problems and solutions," *working paper*. Available at: https://www.jonathandroth.com/assets/files/LogUniqueHOD0_Draft.pdf.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez (2014). "Where is the land of Opportunity? The Geography of Intergenerational Mobility in the United States," *Quarterly Journal of Economics*, 129(4): 1553-1623.

Cinelli, Carlos, Andrew Forney, and Judea Pearl (2022). "A Crash Course in Good and Bad Controls," *Sociological Methods & Research*, Forthcoming.

Correia, Sergio (2016). "Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator," *working paper*. Available at: http://scorreia.com/research/hdfe.pdf

Correia, Sergio, Paulo Guimaraes, and Thomas Zylkin (2020). "Fast Poisson estimation with high dimensional fixed effects," *The Stata Journal*, 20(1): 95-115.

Eames, Taryn (2024). "Taryn versus Taryn (she/her) versus Taryn (they/them): A Field Experiment on Pronoun Disclosure and Hiring Discrimination." Available at SSRH: https://ssrn.com/astract=4694653. Last accessed: March 7, 2024.

Fisher, Robin, Geof Gee, and Adam Looney (2018). "Same-Sex Married Tax Filers After 'Windsor' and 'Obergefell'," *Demography*, 55(4): 1423-1446.

Flores, Andrew R. (2021). "Social Acceptance of LGBTI People in 175 Countries and Locations," UCLA Williams Institute Report. Available at:

https://williamsinstitute.law.ucla.edu/publications/global-acceptance-index-lgbt/. Last accessed January 16, 2024.

Flores, Andrew R., Jody L. Herman, Gary J. Gates, and Taylor N. T. Brown (2016). "How Many Adults Identify as Transgender in the United States," Williams Institute Report. Available at: http://williamsinstitute.law.ucla.edu/research/how-many-adults-identify-as-transgender-in-the-united-states/.

Geijtenbeek, Lydia, and Erik Plug (2018). "Is there a penalty for registered women? Is there a premium or registered men? Evidence from a sample of transsexual workers," *European Economic Review*, 109: 334-347.

Grant, Jaime M., Lisa A. Motter, and Justin Tanis. "Injustice at every turn: A report of the national transgender discrimination survey." (2011).

Haltiwanger, John, Henry Hyatt, and Erika McEntarfer (2018). "Who Moves Up the Job Ladder?" *Journal of Labor Economics*, 36(S1): S301-S336.

James, Sandy E., Jody L. Herman, Laura E. Durso, and Rodrigo Heng-Lehtinen (2024). "Early Insights: A Report of the 2022 U.S. Transgender Survey." Washington, DC: National Center for Transgender Equality.

James, Sandy, Jody Herman, Susan Rankin, Mara Keisling, Lisa Mottet, and Ma'ayan Anafi (2016) "The report of the 2015 US transgender survey." Washington, DC: National Center for Transgender Equality.

Jones, Jeffrey M. (2024). "LGBT Identification in U.S. Now at 7.6%," Gallup News, March 13, 2024. Available at: https://news.gallup.com/poll/611864/lgbtq-identification.aspx.

Kline, Patrick, Evan K. Rose, and Christopher R. Walters (2022). "Systemic Discrimination Among Large U.S. Employers," *The Quarterly Journal of Economics*, 137(4): 1963-2036.

Kolk, Martin, J. Lucas Tilley, Emma von Essen, Ylva Moberg, and Ian Burn (2023). "The Demography of Sweden's Transgender Population – Patterns, Changes, and Sociodemographics," SSRN Working Paper #4427508.

Kunze, Astrid (2017). "The Gender Wage Gap in Developed Countries," in Susan L. Averett, Laura M. Argys, and Saul D. Hoffman (eds.), *The Oxford Handbook of Women and the Economy*.

Kychenthal, Skye (2022). "Why the Trans Programmer?" Working Paper. Available at: https://arxiv.org/pdf/2205.01553.pdf

Mann, Samuel (2021). "Transgender employment and gender marker laws," *Labour Economics*, Article 102072.

McDowell, Alex, Ana M. Progovac, Benjamin Le Cook, and Sherri Rose (2019). "Estimating the Health Status of Privately Insured Gender Minority Children and Adults," *LGBT Health*, 6(6): 289-296.

MIT Election Data and Science Lab (2018). "County Presidential Election Returns 2000-2020," Harvard Dataverse, V11. Available at: https://doi.org/10.7910/DVN/VOQCHQ.

Movement Advancement Project and Center for American Progress (2015). *Paying an Unfair Price: The Financial Penalty for Being Transgender in America.* Available at: https://www.lgbtmap.org/policy-and-issue-analysis/unfair-price-transgender

Movement Advancement Project (2023). "Equality Maps: Bans on Best Practice Medical Care for Transgender Youth." Available at: https://www.mapresearch.org/equality-maps/healthcare/youth_medical_care_bans

National Center for Transgender Equality (2013). "Transgender People and the Social Security Administration." Available at: https://transequality.org/sites/default/files/docs/kyr/SSAResource_June2013.pdf. Last accessed March 24, 2024.

National Center for Transgender Equality (2024). "Know your Rights: Social Security". Available at: https://transequality.org/know-your-rights/social-security#:~:text=How%20is%20the%20Gender%20in,date%20of%20birth%2C%20and%20gender. Last accessed January 15, 2024.

O'Donnabhain v. Commissioner, 134 T.C. 34 No. 4 (Feb 2, 2010), *Doc 2010-2490*, *2010 TNT 22-8*.

Sansone, Dario (2019). "LGBT students: New evidence on demographics and educational outcomes." *Economics of Education Review*, 73: Article 101933.

Schilt, Kristen and Matthew Wiswall (2008). "Before and After: Gender Transitions, Human Capital, and Workplace Experiences," *The B.E. Journal of Economic Analysis & Policy*, 8(1): Article 39.

Shannon, Matthew (2021). "The Labour Market Outcomes of Transgender Individuals," *Labour Economics*, Article 102006.

Stacey, Lawrence, Rin Reczek, and R. Spiker (2022). "Toward a Holistic Demographic Profile of Sexual and Gender Minority Well-being," *Demography*, 59(4): 1403-1430.

Figure 1a: Counts of Administrative First Name and Gender Marker Changes Have Increased Over Time

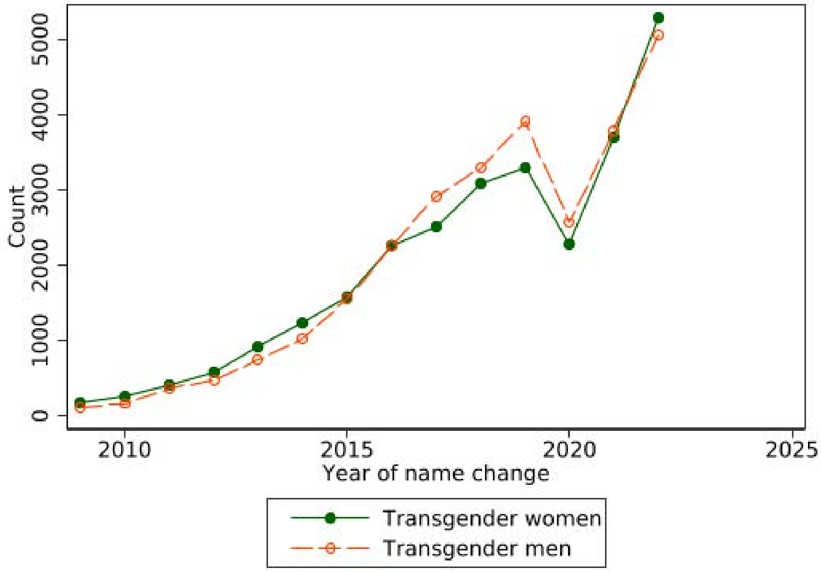

Figure 1b: Transgender Men Transition at Younger Ages than Transgender Women on Average

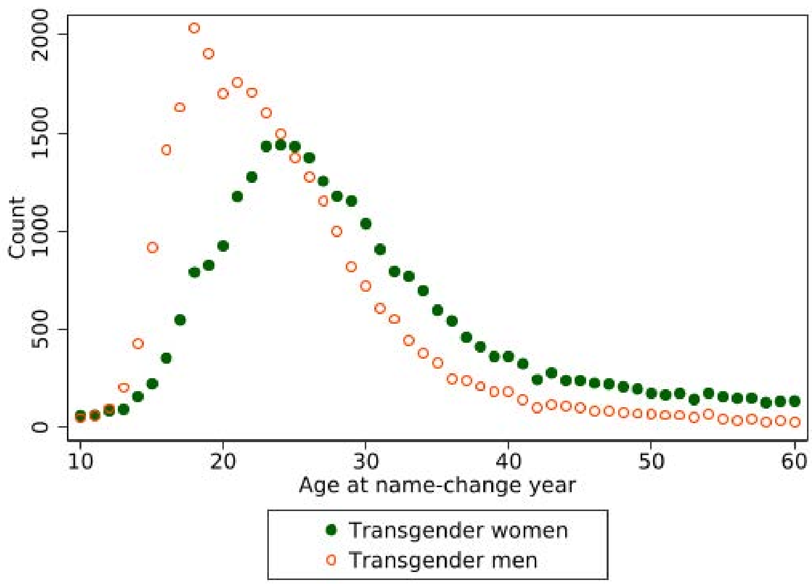

Notes: Author calculations from confidential IRS data. See text for details on how the name-change year is determined.

Figure 2a: Transgender Men (as a Share of State Population) are Concentrated in More Progressive States

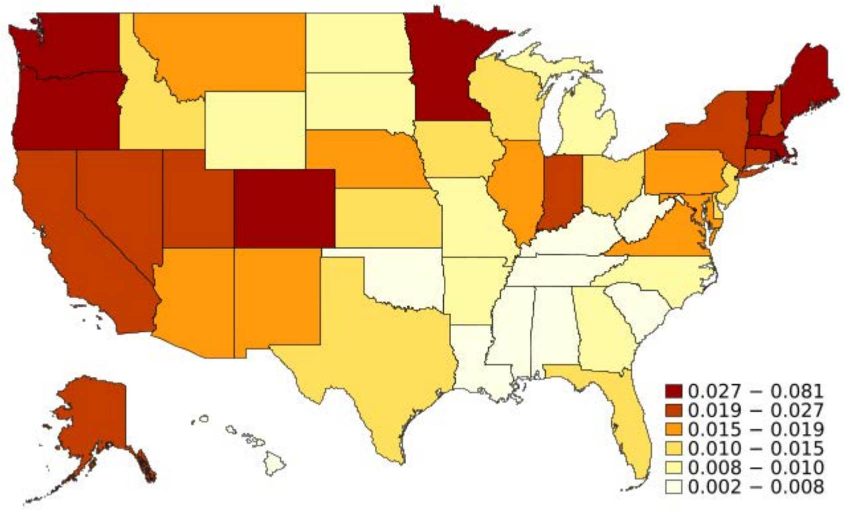

Figure 2b: Transgender Women (as a Share of State Population) are Concentrated in More Progressive States

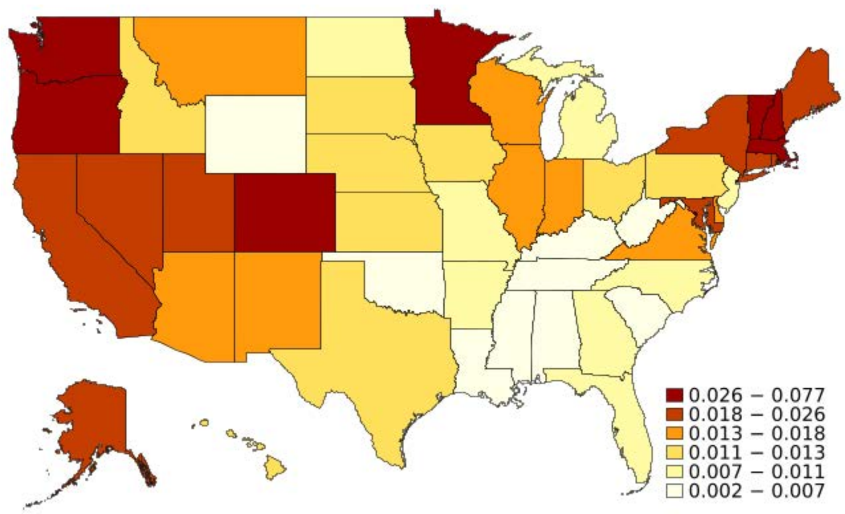

Notes: Author calculations from confidential IRS data. In Figure 2a we calculate the population share as the ratio of transgender men (those who were assigned female at birth and changed their gender marker from female to male) as a share of all individuals assigned female at birth (calculated using the control sample). In Figure 2b we calculate the population share as the ratio of transgender women (those who were assigned male at birth and changed their gender marker from male to female) as a share of all individuals assigned male at birth (calculated using the control sample). The shares reported in the legends are in percent (i.e., multiplied by 100).

Figure 3: Poisson Event Study Estimates of Annual Earnings

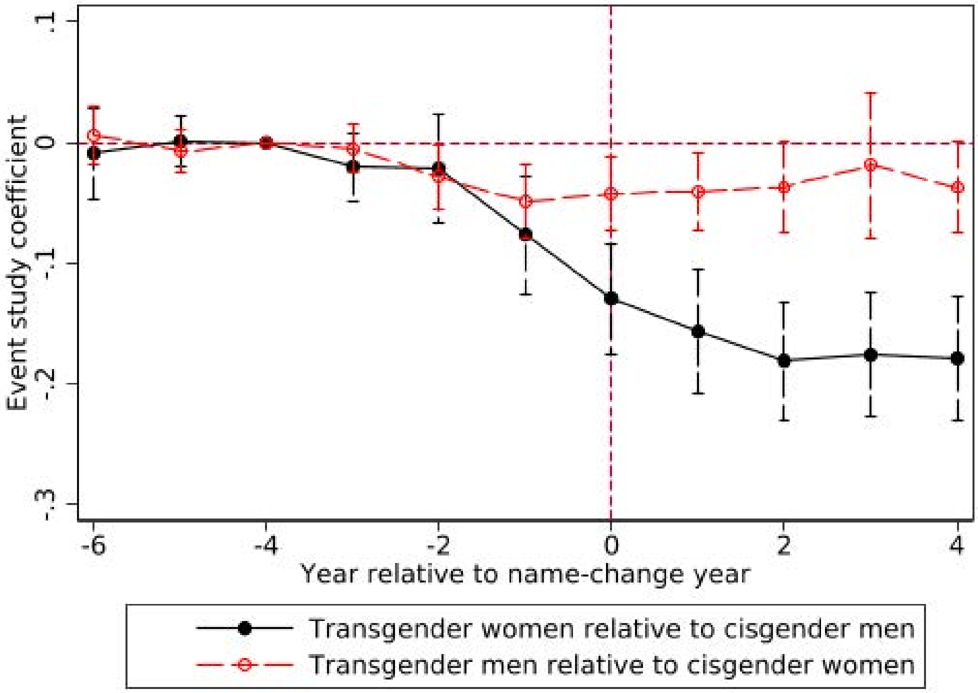

Notes: Author calculations from confidential IRS data. Shown are the coefficients of a Poisson stacked event study regression. The solid black line is the earnings gap for transgender men (those who were assigned female at birth and changed their gender marker from female to male) relative to cisgender women as control. The dashed red line is the earnings gap for transgender women (those who were assigned male at birth and changed their gender marker from male to female) relative to cisgender men as control. Standard errors are clustered by unique individual. Each sample (transgender women or men and cisgender men or women) is reweighted to match the age distribution of transgender individuals (aggregated between transgender men and transgender women). Event time -4 is omitted. We restrict to transgender individuals who are age 28 or later in the name-change year and whose name-change year is 2018 or earlier. See text for further details of the stacked event study methodology.

Figure 4: Event Study Estimates of Extensive Margin Participation

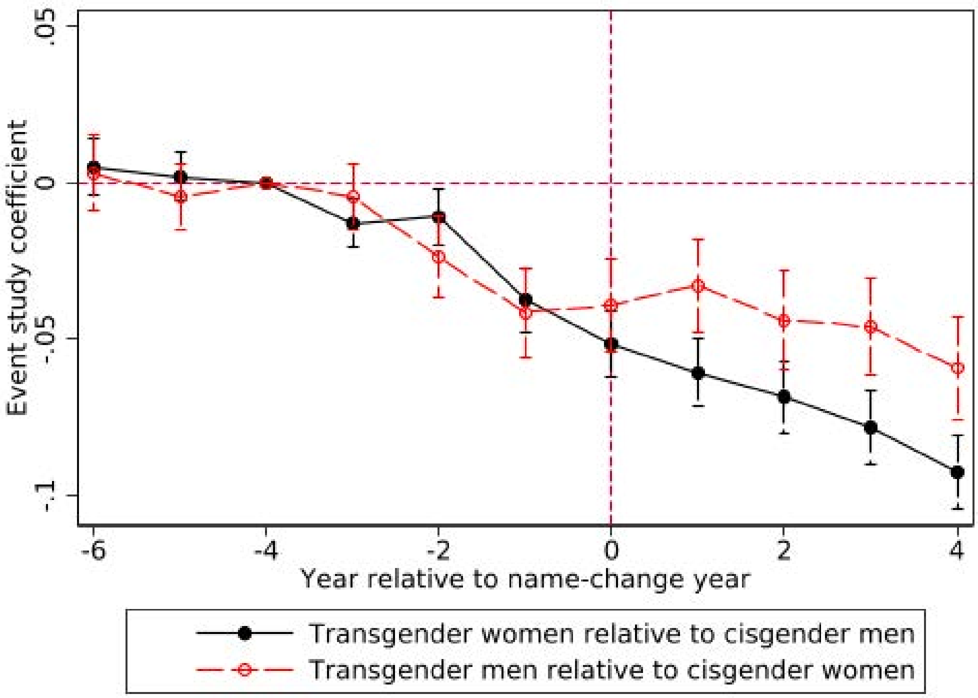

Notes: Author calculations from confidential IRS data. Shown are the coefficients of a stacked event study regression. The solid black line is the change in labor market participation for transgender men (those who were assigned female at birth and changed their gender marker from female to male) relative to cisgender women as control. The dashed red line is the change in labor market participation for transgender women (those who were assigned male at birth and changed their gender marker from male to female) relative to cisgender men as control. Each sample (transgender women or men and cisgender men or women) is reweighted to match the age distribution of transgender individuals (aggregated between transgender men and transgender women). Event time -4 is omitted. We restrict to transgender individuals who are age 28 or later in the name-change year and whose name-change year is 2018 or earlier. See text for further details of stacked event study methodology.

Table 1: Descriptive Statistics on Age, Geography, Education, Family, and Childhood Background Outcomes

	(1)	(2)	(3)	(4)
Variable	Transgender men	Cisgender men	Transgender women	Cisgender women
From the individual's own tax records:				
Age	28.1	28.1	34.6	34.6
County share of votes for D. in 2000-20 Pres. Elections	0.585	0.525	0.601	0.530
% Black in zip code from 2013-17	0.100	0.126	0.105	0.130
% White in zip code from 2013-17	0.737	0.728	0.723	0.726
Share with any college	0.875	0.643	0.876	0.758
Number of years of college	5.071	3.028	4.969	4.102
Married	0.207	0.288	0.212	0.429
Married to someone we identify as transgender, among married	0.044	0.000	0.060	0.000
Any dependents	0.084	0.250	0.081	0.451
N	28,230	353,490	27,580	357,520
From parent/child tax linkages:				
% Individuals linked to a parent (1977 or later cohorts)	0.947	0.897	0.949	0.876
Parents' income (in 2022 dollars)	116,301	85,668	135,505	85,112
Single parent household	0.235	0.327	0.204	0.324
% White in parent's zip code (2000)	0.792	0.749	0.798	0.746
N	26,700	245,470	22,950	242,670

Author calculations of sample means from confidential IRS data. The sample of cisgender men (women) is a random sample stratified by cohort of those who filed a 2022 tax return, appeared as a dependent on a 2022 tax return, or had an information return in 2022. The marital status and dependents means are conditioned on being a tax filer. The means for parental income, single parent household, and % white in parent's zip code are conditioned on being linked to a parent. The sample of cisgender men (women) has been re-weighted to match the year and cohort distribution of transgender men (women), separately for each row. The sample size in the bottom panel represents the counts of individuals in 1977 or later cohorts. For disclosure-avoidance purposes, we round sample counts to the nearest multiple of 10.

Table 2: Siblings Fixed Effects Estimates of Transgender Earnings Gaps

	No controls (Sibling pair fixed effects only)	Education and potential experience controls	Full controls
	Transgender women and cisgender brother sibling pairs		
Transgender	-0.173 (0.023)	-0.192 (0.022)	-0.217 (0.021)
	Transgender men and cisgender sister sibling pairs		
Transgender	0.017 (0.018)	0.019 (0.017)	-0.038 (0.018)
	All transgender and cisgender sibling pairs		
Transgender	-0.079 (0.011)	-0.083 (0.011)	-0.130 (0.010)
Male	0.095 (0.015)	0.112 (0.015)	0.133 (0.013)
Male at birth	0.202 (0.015)	0.186 (0.015)	0.062 (0.015)
Trans N	20,870	20,870	20,870

Author calculations from confidential IRS data. This table reports estimates using a sample of transgender individuals linked to their siblings, as discussed in Section IV.B. We restrict to transgender individuals and siblings within the 1981-1999 birth cohorts. “Education and potential experience controls” includes a linear term for number of years of education, a dummy for currently being a college student, a dummy for currently being enrolled more than half time, and a quadratic in potential experience. “Full controls” adds a dummy for being married, a dummy for having a youngest dependent under age 6, and a dummy for having a youngest dependent age 6 and above, fixed effects for two-digit NAICS code of the highest-wage employer (including missing), fixed effects for two-digit occupation (including missing), log population density of the zip code, and a dummy for being in a “central county” in one of the top 30 MSAs ranked by population. Standard errors are clustered by unique transgender individuals. For disclosure-avoidance purposes, we round sample counts to the nearest multiple of 10.

Table 3: Coworker Fixed Effects Estimates of Transgender Earnings Gaps

	No controls (Coworker pair fixed effects only)	Education and potential experience controls	Full controls
	Transgender women and cisgender men coworker pairs		
Transgender	-0.109 (0.010)	-0.104 (0.010)	-0.083 (0.010)
	Transgender men and cisgender women coworker pairs		
Transgender	-0.034 (0.009)	-0.031 (0.009)	-0.050 (0.009)
	All transgender and cisgender coworker pairs		
Transgender	-0.086 (0.007)	-0.081 (0.007)	-0.076 (0.007)
Male	0.033 (0.007)	0.038 (0.007)	0.029 (0.007)
Male at birth	0.060 (0.011)	0.054 (0.011)	0.055 (0.010)
Trans N	13,230	13,230	13,230

Author calculations from confidential IRS data. This table reports estimates using a sample of transgender individuals linked to their coworkers, as discussed in Section IV.C. We restrict to transgender individuals and coworkers in same three-digit occupation working for same employer within the 1981-1999 birth cohorts. “Education and potential experience controls” includes a linear term for number of years of education, a dummy for currently being a college student, a dummy for currently being enrolled more than half time, and a quadratic in potential experience. “Full controls” adds a dummy for being married, a dummy for having a youngest dependent under age 6, and a dummy for having a youngest dependent age 6 and above, log population density of the zip code and a dummy for being in a “central county” in one of the top 30 MSAs ranked by population. For disclosure-avoidance purposes, we round sample counts to the nearest multiple of 10.

Online Appendix – Not for Publication

Transgender Earnings Gaps in the United States: Evidence from Administrative Data

A. Identifying likely transgender individuals and labor market outcomes

We identify likely transgender individuals using information on gender marker changes from the Social Security Administration (SSA) matched to gender-consistent name changes from tax records at the Internal Revenue Service. The Social Security Administration maintains records of all individuals identified by a Social Security Number (SSN), including date of birth (and death, if applicable) and gender marker. The gender marker is not used to administer SSA programs but is used in some circumstances for identity verification or to disallow certain publicly-funded medical procedures that are incongruent with that SSA gender marker (National Center for Transgender Equality 2024). The SSA gender marker is binary; there is no non-binary option. The IRS receives this data from the SSA and assembles a dataset with one observation per person, reflecting the most up-to-date record for each individual. We refer to this dataset as the "SSA Gender Marker Dataset." We have access to the current vintage of the SSA Gender Marker Dataset (which we downloaded in May 2024), as well as a vintage from 2011.[34] Thus, we can observe if a given SSN has a different gender marker currently than it did in the 2011 vintage.[35] In total, we find approximately 305,900 individuals whose gender marker changed in the SSA database.

However, it would be misleading to rely on SSA gender marker changes alone. A fundamental challenge is that (a) there exist clerical errors in the SSA data and (b) the base rate of being transgender is relatively low. To be precise,

[34] The version of the data from the SSA that we have access to contains SSNs but not first names. We observe first names in the IRS tax records. We remind the reader that all data work was performed by Goodman on IRS servers.

[35] We restrict to those observations where the date of birth matches between the 2011 and current vintages to mitigate against the possibility that the record reflects two different people.

from Bayes' rule, the probability that an individual is transgender given a gender marker change can be given as follows, where gm denotes a gender marker change:

$$\Pr(trans|gm) = \frac{\Pr(gm|trans)\Pr(trans)}{\Pr(gm|trans)\Pr(trans) + \Pr(gm|cis)\Pr(cis)} \qquad \text{(A1)}$$

While intrinsically $\Pr(gm|trans) \gg \Pr(gm|cis)$, it is also the case that $\Pr(cis) \gg \Pr(trans)$, so the second term in the denominator of (A1) might not be small relative to the first term, causing $\Pr(trans|gm)$ to be substantially less than one. As a result, such a naïve approach would lead to substantial false positives.

For this reason, in our methodology, we identify individuals as likely transgender only if we also observe a first name change in the 'gender congruent direction' in the tax records, in the spirit of Fisher, Gee, and Looney (2018). That is, for an individual whose gender marker changes from 'M' to 'F', we require that the first name changes to a name that is more female, and vice versa. Let nc denote such a correct-direction name change. The probability that an individual is transgender given a gender marker change and a correct-direction first name change is:

$$\Pr(trans|nc\&gm) = \frac{\Pr(nc\&gm|trans)\Pr(trans)}{\Pr(nc\&gm|trans)\Pr(trans) + \Pr(nc\&gm|cis)\Pr(cis)} \qquad \text{(A2)}$$

This additional requirement substantially improves the specificity (i.e., one minus the rate of false positives) under the assumption that $\frac{\Pr(nc\&gm|trans)}{\Pr(nc\&gm|cis)} \gg \frac{\Pr(gm|trans)}{\Pr(gm|cis)}$. That is, while cisgender individuals do occasionally correct clerical gender marker errors (so that $\Pr(gm|cis)$ is non-trivial), it is exceedingly rare for cisgender individuals to *also* change their first names in a manner consistent with a gender transition.

To identify name changes, we use the first names written on IRS Form 1040, which we observe from 2008 through 2022. The instructions to Form 1040 ask taxpayers to write the full legal names of the primary filer, secondary filer, and any dependents on Form 1040; we link each of these individuals to the SSA Gender Marker Dataset via SSN.[36] We drop any name changes that may be spurious. These include situations where: (a) the name is a single initial, (b) the name is identical to a spouse's name in a nearby year, (c) the name is the same as a middle or last name in some other year, (d) the name is identical to the combination of the first and middle name in some other year, (e) the new name is a diminutive of the old name, or vice versa (such as "CHRIS" and "CHRISTOPHER"), or (f) in the case of a dependent, the name is the same as the name of a sibling claimed in some other year. Finally, we restrict attention to observations where we observe exactly one name change, so that we can cleanly identify a "pre-transition" name and a "post-transition" name.

Next, we use the SSA's published database of baby names to compute the gender share of each name. This database reports the number of births with a given name j in a given year c', along with the share of those babies who were assigned male at birth. In our procedure, we assign a gender share to a given observation – defined by the name and the year of birth c – by aggregating over all cohorts within five years of the focal individual (i.e., $|c - c'| \leq 5$). Then, for each individual, we compute the change in gender congruence between their two names. For example, for an individual who changes gender marker from "M" to "F", the change in gender congruence is equal to the share female of the post-transition name minus the share female of the pre-transition name. We also

[36] In cases when a given SSN is both a filer (primary or secondary) and a dependent in the same year, we discard the dependent observation.

compute the confidence interval of that change, taking into account the "sampling" error in the SSA baby name data – that is, we are less sure of the change in gender congruence when one or both of the names is uncommon.[37] We consider a name change to be "gender congruent" when the bottom of the confidence interval of the gender congruence is greater than or equal to twenty percentage points. When we apply all of these selection rules, we are left with approximately 55,560 likely transgender individuals, of whom 28,230 changed their gender marker to "male" and switched to a more masculine name (i.e., transgender men) and 27,580 changed their gender marker to "female" and switched to a more feminine name (i.e., transgender women).

The change in the sample size from each rule we apply is described in Appendix Table A1. The vast majority of those we drop fail to exhibit any name change; of the 250,100 people we drop in total, approximately 200,000 people do not change their first name on Form 1040. An additional 16,000 never appear on any Form 1040 beginning in 2008. We drop approximately 13,500 individuals with a gender marker change and a single valid name change due to their pre-change or post-change name being too uncommon to be found in the SSA baby names database (about 8,070) or their name change being insufficiently gendered (about 5,430). We also drop about 340 people whose name change occurs prior to age 10. Our sample construction is not sensitive to the choice to use a 20 percentage point threshold for determining whether a name change is sufficiently gendered; among those 55,800 individuals in our sample, 49,200 would have survived an 80 percentage point threshold rather than a 20 percentage point threshold. Additionally, among the 5,430 who fail to achieve the 20 percentage

[37] Formally, we treat the change in the gender congruence as a change in binomial proportions, taking the "successes" and "trials" from the SSA baby name data. To compute this confidence interval, we use a simple finite-sample adjustment proposed by Agresti & Caffo (2000).

point threshold, 3,810 have essentially no change in gender congruence (less than 5 percentage points in magnitude).

It is important to be explicit about the limitations of our approach for identifying transgender people in administrative data. First and most importantly, not all transgender people will change their gender marker and/or name on federal documents. Thus, the samples of likely transgender people we identify are smaller than the sample sizes implied by recent nationally representative surveys.[38] It is likely that the people we identify have taken more steps with respect to gender affirmation than people who do not pursue these types of changes. It is also likely that people we identify are positively selected considering the burdensome process of changing gender marker and first name on federal documents. Second, our name change algorithm will miss people whose birth or chosen name is very uncommon, or whose name change is not sufficiently gendered. To take a specific hypothetical example, an individual born in 1990 who changed their name from Kerry to Casey would not be identified as likely transgender, as this name change does not exceed the twenty percentage point threshold.[39]

[38] Our sample size is in the same order of magnitude as the samples identified in U.S. health claims data. McDowell et al. (2019) estimate a transgender share in their population of 0.026%, compared to 0.017% in our data. Baker and Restar (2022) identify 16,619 transgender individuals in their database of commercially-insured individuals, compared to our sample of 55,800.
[39] This example – which is fictional – uses information solely from the publicly-available SSA baby name database.

Appendix Table A1: Counts of SSA Gender Marker Changes

	N
Total with gender marker change	305,900
Drop if never appear on a tax return	289,870
Drop if literal first name string is constant	90,040
Drop if no valid name changes	80,040
Drop if multiple valid name changes	69,640
Drop if name not found in baby name database	61,570
Drop if name change occurs prior to age 10	61,230
Drop if gender share change is less than 0.2	55,800

Author calculations from confidential IRS data. Invalid name changes are the use of diminutives, initials, and name transposed with sibling or spouse. For disclosure-avoidance purposes, we round sample counts to the nearest multiple of ten.

B. Additional validation of algorithm

We provide additional validation on our name-change approach to identifying likely transgender individuals. First, we use an auxiliary SSA-served dataset. Specifically, in addition to the main SSA Gender Marker Dataset, which lists the most updated record for each individual presently and in 2011, we have access to an auxiliary dataset that lists all updates to an individual's SSA record beginning in 2015. We refer to this latter dataset as the "Record Update Dataset". Relative to the SSA Gender Marker Dataset, the Record Update Dataset has two main advantages: (1) it includes first names and (2) it specifies the exact date (defined at the weekly level) of the record update. However, it also has two main disadvantages: (1) it begins in 2015 – later than the 2011 vintage of the SSA Gender Marker Dataset that we have access to – and (2) it does not indicate what the record was changed *from*. As a fictional example, the record may indicate an update on April 8, 2018 for a person with name Jane Doe, SSN XXX-XX-0001, and gender marker "F". But, unless this same person made a record update at some point between January 1, 2015 and April 8, 2018, the dataset does not indicate whether this update reflects a name change, a gender marker change, or a change of something else altogether (e.g., citizenship).

For these reasons, the Record Update Dataset is not suitable for use as the main mechanism to identify transgender individuals. However, we use it for validation in two ways. First, we can use it to validate the name changes in Form 1040 from 2015 onward. Because taxpayers are instructed to use their SSA names on their Form 1040, we expect to see SSA record updates for these individuals with a first name matching what we see on Form 1040. Second, we can use the Record Update Dataset to infer the likely timing of gender marker and, if applicable, name changes (as recorded by the SSA) for most individuals with a gender marker change. Specifically, we can identify the date of the first record with the post-transition name and/or post-transition gender marker. We note that this will be inaccurate if the name change or gender marker change occurred prior to 2015 and if a subsequent record update of any kind is made in 2015 or later, but given the evidence in Figure 1 in the main text regarding the timing of name changes on Form 1040, we expect this concern to be second order.

Appendix Table B1 presents various statistics for four sets of individuals with a gender marker change: those with no name change at all (column 1), those with a name change but whose pre-change and/or post-change names are not found in the SSA Baby Names dataset (column 2), those with a name change that is not congruent with the gender marker change (column 3), and our sample (column 4). In the first set of rows, we provide evidence that a substantial minority of name changes in columns 2 and 3 reflect minor spelling changes or corrections on Form 1040, rather than an actual change in one's legal name. First, we find that the average Levenshtein string distance between the two names is much smaller in these columns. Second, we compare the Form 1040 post-transition names (restricted to those with name-change year in 2015 or later) to names recorded in the Record Update Dataset; we find that 98% of our sample (column 4) have an SSA record update with a name matching the Form 1040 post-transition name, while this share is approximately 75-80% for columns 2 and 3.

The remainder of Appendix Table B1 presents additional statistics for these groups. Those with a name change involving an uncommon name, or with a name change that is not congruent with the gender marker change, have observables that tend to be between those without a name change and those we identify as transgender. Those with a gender marker change who do not make our sample are much less likely to be married to another individual we identify as transgender.

Next, we examine how well the Form 1040 timing of name changes matches the timing of name change in the Record Update Dataset. Each row in

Appendix Table B2 corresponds to a year of name change according to Form 1040. Each column reports the conditional probability of the Record Update Dataset name change being in a given year; the rows each sum to one. Appendix Table B3 reports the analogous figures for the timing of gender marker changes in the Record Update Dataset. In both cases, for those with Form 1040 name change year equal to t, approximately 80% have a record update change in t or $t+1$. Recall that tax returns for year t are typically filed in year $t+1$; thus, it is unsurprising that many year t tax returns reflect names updated in year $t+1$.

In Appendix Table B4, we report the apparent timing of gender marker changes relative to name changes. We estimate that, of those individuals for whom we can ascertain the timing of both changes, 89% change both their name and their gender marker in the same month.[40] 5.3% change their gender marker 1-12 months after their name change and 3.8% do so 13 or more months after their name change. Approximately 1.1% change their gender marker prior to their name change.

Finally, Appendix Figure B1 plots the apparent timing of gender marker change based on when the first record with the updated gender marker is observed. We do so both for the baseline dataset, as well as the set of individuals with a gender marker change but no name change (i.e., those in column (1) of Appendix Table B1). Several patterns emerge from this figure. First, among those with no name changes, the prevalence of gender marker changes is relatively constant over time until the beginning of the pandemic, which contrasts with the increasing prevalence over time within our main transgender sample. Second, the number of those with a gender marker change but no (measured) name change increases sharply in 2023, with a corresponding decrease in those with a gender marker change and name change. These may reflect individuals who have in fact

[40] Of these, 99% change both their name and their gender marker in the same week – likely in the same visit to an SSA office.

changed their name, but this name was not reflected on tax returns during our sample period (which ends in 2022).

Appendix Table B1: Descriptive Statistics, Individuals with Gender Marker Change but Without Gender-Congruent Name Change

	(1)	(2)	(3)	(4)
Variable	Literally no name change	Name is rare	Gender share change not >.2	Our sample
Age	50.2	41.1	41.3	31.3
Levenshtein string distance between names	N/A	4.10	3.26	5.14
Share with string distance ≤ 2	N/A	0.324	0.480	0.067
Share with name match in Record Update dataset	N/A	0.810	0.742	0.982
County share of votes for D. in 2000-20 Pres. elections	0.544	0.596	0.566	0.593
% Black in zip code from 2013-17	0.141	0.143	0.142	0.103
% White in zip code from 2013-17	0.700	0.665	0.686	0.731
Share with any college	0.699	0.791	0.757	0.875
Number of years of college	3.558	4.265	40.020	5.019
Married	0.508	0.379	0.351	0.209
Married to someone we identify as transgender	0.001	0.012	0.014	0.052
Any dependents	0.273	0.195	0.228	0.082
N	199,385	8,860	5,430	55,800

Author calculations of sample means from confidential IRS data in 2022. Column 4 refers to the transgender individuals in our final sample. Column 1 includes those whose first name string on Form 1040 is constant. Column 2 includes individuals with exactly one first name change, but whose pre- or post-change names are too uncommon to be included in the SSA baby name database. Column 3 includes individuals whose first name change does not exceed the 0.2 threshold for congruency with the gender marker change. The “married”, “any dependents”, and “number of dependents” rows are conditioned on being a tax filer. The “married to someone we identify as transgender” row is conditioned on being married. Samples have not been reweighted by age. For disclosure-avoidance purposes, we round sample counts to the nearest multiple of 10.

Appendix Table B2: Form 1040 name change year versus Record Update dataset name change year

Year of 1040 name change	No valid record update date	2015 record update	2016 record update	2017 record update	2018 record update	2019 record update	2020 record update	2021 record update	2022 record update	2023 record update	2024 record update
2015	0.073	**0.625**	0.189	0.051	0.023	0.021	---	0.006	0.004	0.004	---
2016	0.064	0.063	**0.640**	0.147	0.040	0.022	0.007	0.005	0.006	---	---
2017	0.033	0.018	0.084	**0.662**	0.151	0.031	0.006	0.006	0.005	---	---
2018	0.017	0.012	0.027	0.088	**0.655**	0.166	0.015	0.006	0.007	0.006	0.002
2019	0.014	0.007	0.015	0.027	0.091	**0.700**	0.105	0.018	0.014	0.007	0.002
2020	0.016	0.003	0.005	0.013	0.029	0.127	**0.517**	0.226	0.046	0.016	0.002
2021	0.012	---	---	0.003	0.008	0.017	0.048	**0.665**	0.203	0.035	0.005
2022	0.013	---	---	0.002	0.002	0.005	0.007	0.050	**0.726**	0.189	0.006

Author calculations from confidential IRS data. Each individual in our sample is defined by a name-change year (rows) and a year, if any, when we observe the first record with the post-transition name in the Record Update dataset (columns). Each cell gives the conditional probability of being in that column, conditional on being in that row. Some cells are suppressed for disclosure avoidance.

Appendix Table B3: Form 1040 name change year versus Record Update gender marker change year

Year of 1040 name change	No valid record update date	2015 record update	2016 record update	2017 record update	2018 record update	2019 record update	2020 record update	2021 record update	2022 record update	2023 record update	2024 record update
2015	0.072	**0.497**	0.218	0.082	0.042	0.038	0.010	0.010	0.012	0.015	0.004
2016	0.061	0.055	**0.557**	0.180	0.061	0.035	0.011	0.010	0.012	0.015	0.003
2017	0.025	0.016	0.080	**0.607**	0.170	0.052	0.009	0.010	0.015	0.013	0.003
2018	0.009	0.011	0.027	0.083	**0.616**	0.184	0.022	0.012	0.016	0.018	0.004
2019	0.004	0.006	0.015	0.027	0.088	**0.668**	0.115	0.026	0.025	0.021	0.005
2020	0.005	0.003	0.006	0.013	0.030	0.126	**0.484**	0.229	0.058	0.039	0.007
2021	**---**	**---**	0.003	0.004	0.009	0.018	0.046	**0.633**	0.215	0.061	0.009
2022	**---**	**---**	0.002	0.002	0.003	0.006	0.007	0.048	**0.709**	0.213	0.010

Author calculations from confidential IRS data. Each individual in our sample is defined by a name-change year (rows) and a year, if any, when we observe the first record with the post-transition gender marker in the Record Update dataset (columns). Each cell gives the conditional probability of being in that column, conditional on being in that row. Some cells are suppressed for disclosure avoidance

Appendix Table B4: Timing of gender marker change relative to name change

Gender marker change relative to name change	Density
More than 12 months before	0.004
1 to 12 months before	0.007
In same month	0.889
1 to 6 months later	0.037
7 to 12 months later	0.016
13 to 24 months later	0.019
More than 24 months later	0.028
N	51,190

Author calculations from confidential IRS data. This table restricts to those where we can observe an apparent name change date and an apparent gender marker date.

Appendix Figure B1: Apparent Timing of Gender Marker Changes

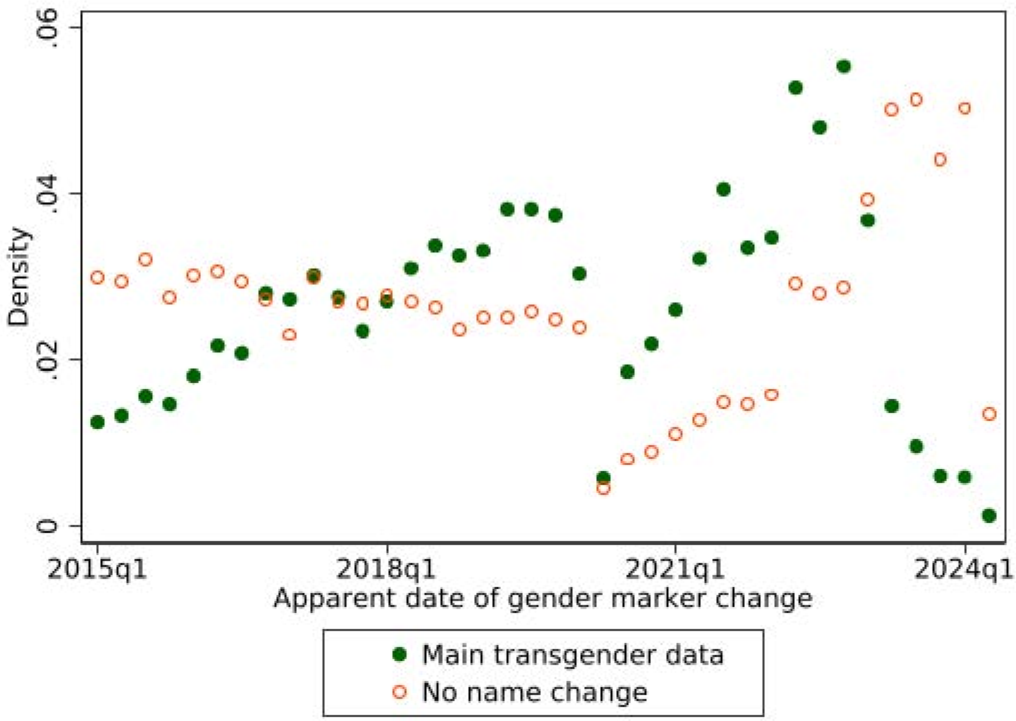

Notes: Author calculations from confidential IRS data. This figure plots the share of observations by apparent quarter of their gender marker change (defined as the first quarter when a record update appears with their current gender marker in the Record Update dataset). The solid circles plot the density for our main data of transgender individuals. The hollow circles plot the density for those with a gender marker change but no name change (i.e., those dropped between rows 2 and 3 of Appendix Table A1).

C. Construction of control sample

We identify the population of all individuals with SSNs that (a) appear as a primary filer, secondary filer, or dependent on a 2022 tax return or (b) receive an information return, such as Form W-2 (for those receiving wage income) or 1099-SSA (for those receiving Social Security income). We require the tax return or information return to indicate an address in the 50 states or the District of Columbia. We then take a stratified random sample of this population. The sampling rate is 0.4% for the 1970-1999 cohorts and 0.1% for all other cohorts. We primarily use this sample after reweighting it to match the cohort distribution of transgender individuals; such oversampling increases power given the actual cohort distribution of transgender individuals.

D. Additional detail regarding matching to parents and parents' income measurement

We observe dependent linkages from 1994 onward. Following Chetty et al. (2014), we assign each individual to the first parent (or set of two parents) that claim them as a dependent. Again following Chetty et al. (2014), we impose an age requirement to minimize cases of grandparents claiming a child as a dependent. In particular, for a dependent link to be "valid", we require either (a) the presence of a female claimant between the ages of 15-40 or (b) the presence of a male claimant between the ages of 15-40 and the absence of a female claimant. If a given year involves an "invalid" dependent link, we iterate forward in time until we find a valid one. We then compute the average income of the parents (at the household level) over a five-year period, using Form 1040. For cohorts younger than 1990, we take the average from 1996-2000, the first five years of Form 1040 income data available to us. For cohorts 1990 and older, we take the average of the years when the child was ages 6 to 10.

We measure income at the household (technically, tax unit) level. If an individual is matched to a single parent who subsequently gets married, income includes the income of the original parent as well as the (new) "step-parent" after that marriage takes place. If an individual is matched to married parents who subsequently divorce, we take the sum of their individual incomes from their newly separate Forms 1040. If one of those divorced parents subsequently remarries, we assign half of the new couple's income to the original parent. We define the parents' zip code as the most common zip code reported on Form 1040 in those same years.

E. Evidence on medical expense deductions, moves, and marital status changes

In the main text, we provide summary statistics on name-change timing and age at name-change. Here, we provide additional evidence on the timing of medical expense deductions, geographic moves, and marital status changes to provide further evidence that our dataset does not contain a substantial number of false positives. Appendix Figures E1a and E1b provide descriptive evidence regarding medical expense deductions claimed by individuals in our main sample. In general, taxpayers can claim medical expenses as an itemized deduction to the extent that those expenses exceed 7.5% or 10% (depending on the year) of adjusted gross income. Since a 2011 court case, the IRS has proclaimed that health care related to gender affirmation (including hormone therapy and gender affirmation surgery) is an eligible medical expense for the purpose of this deduction.[41] Such health care can involve substantial costs – which may or may not be covered by insurance – that can plausibly exceed 7.5% or 10% of adjusted gross income. For example, Baker and Restar (2022) estimate an average total cost in the range of $45,000-$65,000 for phalloplasty and vaginoplasty procedures ("bottom surgery") and $15,000-$20,000 for mastectomy and mammoplasty procedures ("top surgery").

We present statistics on medical expense deductions as raw means by "event time", where time zero is the name-change year. Specifically, we plot the likelihood of claiming any medical expenses as an itemized deduction (Figure E1a) and the mean value of those deductions, including zeroes (Figure E1b). For comparability between the samples of transgender women and transgender men, we reweight each sample to share the same distribution of name-change year and year of birth. (We perform this reweighting throughout this appendix.) While we stress that social and medical transitions are unlikely to have perfectly coincided

[41] O'Donnabhain v. Commissioner, 134 T.C. 34 (2010).

with the timing of the transgender individual's name change (and in most cases we suspect they started several years prior to the name change), it is interesting to see in Figures E1a and E1b that medical expense deductions increase sharply coincident with the timing of the change in the individual's first name. Moreover, the pattern of mean medical expense deductions across both samples matches what is known about the nature and frequency of medical procedures that are most common for each type of gender affirmation. In particular, Baker and Restar (2022) find that the costliest procedure – bottom surgery – is nearly four times as common for transgender women than transgender men.

Appendix Figures E2a and E2b provide evidence on the likelihood an individual moves across state lines and the Democratic vote share of the individual's zip code, respectively. Appendix Figure E2a shows that there is a noticeable jump in the probability that transgender men and transgender women move across state lines around the time of their name change. Appendix Figure E2b shows that transgender men and transgender women are both moving to areas with higher Democratic vote shares over time, but that this pattern is more stark for transgender women than transgender men.

Appendix Figures E3 shows the probability that individual is married. The probability of marriage is increasing for transgender men (consistent with general patterns of aging into adulthood), while it is flat or slightly declining for transgender women. Appendix Figures E4a and E4b show entry into marriage and exit from marriage via divorce.[42] The figures show increases in marriage entries and exits, but the entry effect is larger for transgender men and the exit effect is much larger for transgender women.

[42] Exit from marriage due to the death of a spouse is very rare for this sample.

Appendix Figure E1a: Likelihood of Medical Expense Deduction Increases Sharply Coincident with Timing of Name Change Among Transgender Sample

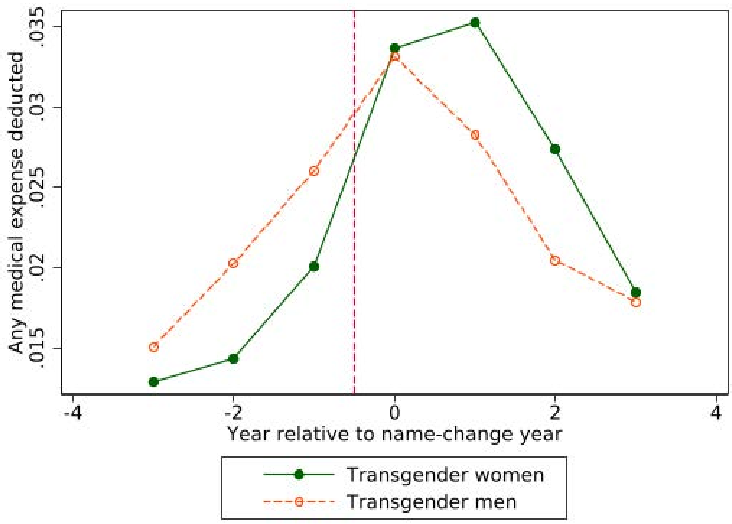

Appendix Figure E1b: Average Medical Expense Deduction Increases Sharply Coincident with Timing of Name Change Among Transgender Sample

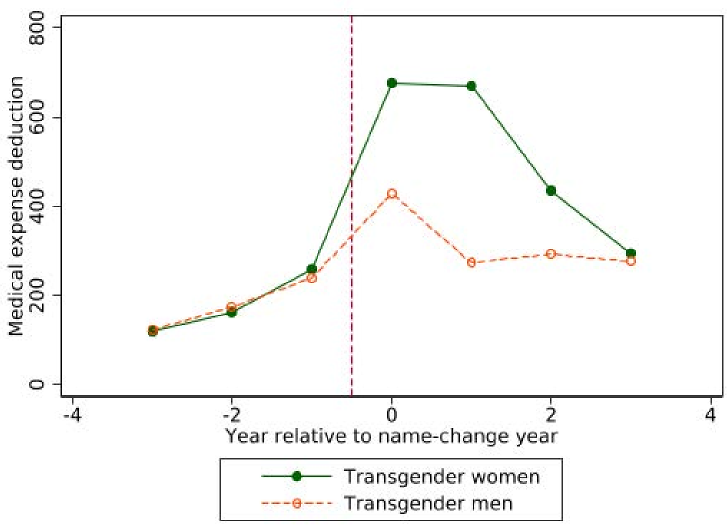

Notes: Author calculations from confidential IRS data. Each sample is reweighted to match the average distribution by age and name-change year. We restrict to individuals whose name-change year is 2019 or earlier. Means in Figure E1b include the zeros.

Appendix Figure E2a: Cross state moves

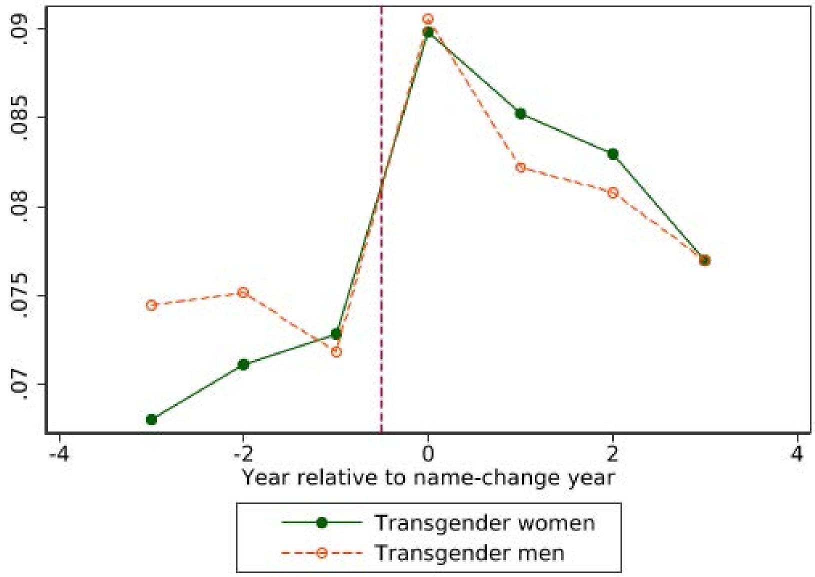

Appendix Figure E2b: Democratic vote share

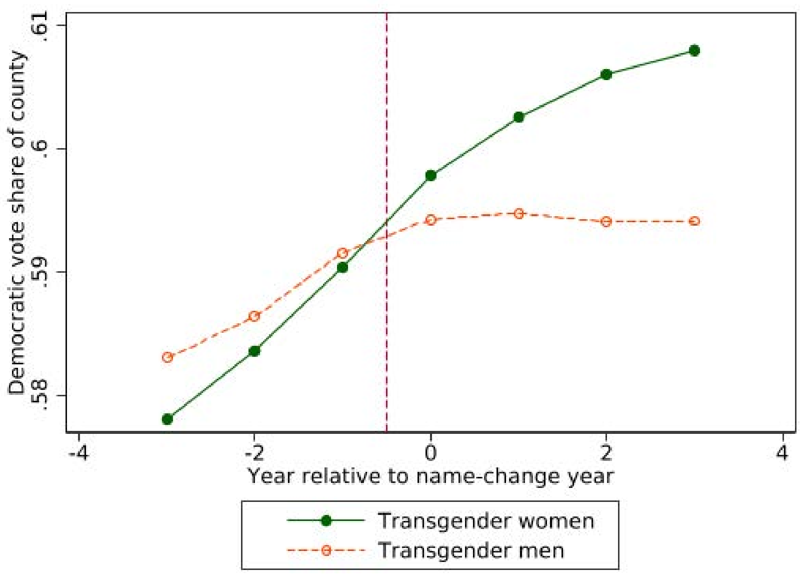

Notes: Author calculations from confidential IRS data. Each sample is reweighted to match the average distribution by age and name-change year. We restrict to individuals whose name-change year is 2019 or earlier. The Democratic vote share is measured at the county level; it is defined as the share of votes received by the Democratic presidential candidate from 2008 through 2020, relative to the total number of votes received by the Democratic and Republican candidates.

Appendix Figure E3: Married

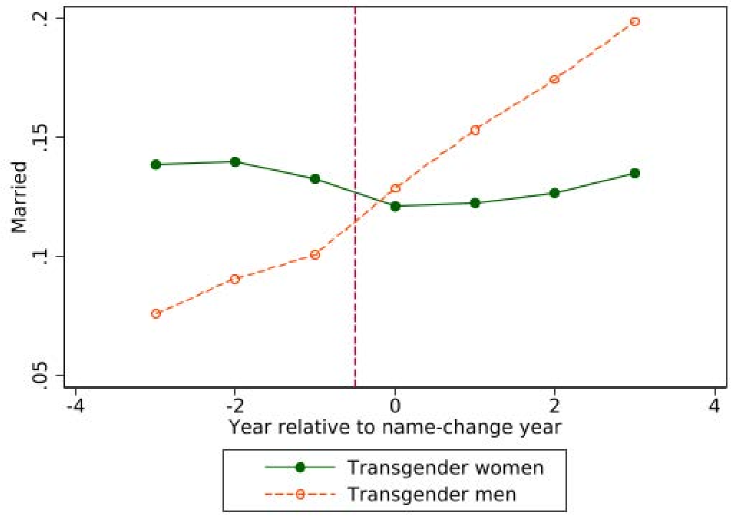

Notes: Author calculations from confidential IRS data. Each sample is reweighted to match the average distribution by age and name-change year. We restrict to individuals whose name-change year is 2019 or earlier.

Appendix Figure E4a: Entry into marriage

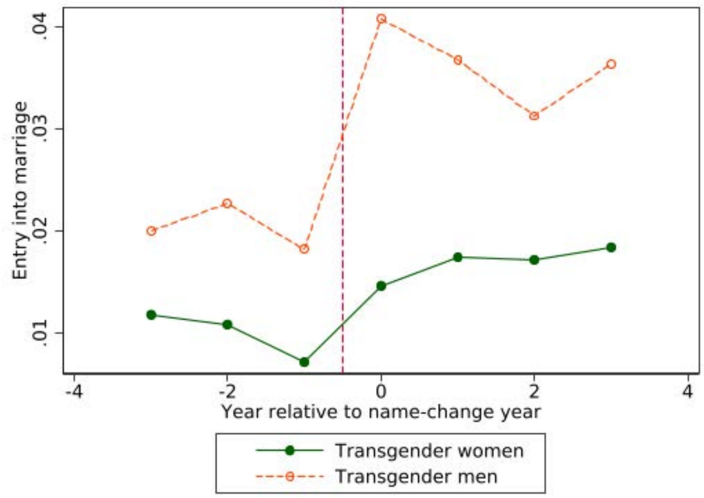

Appendix Figure E4b: Divorce

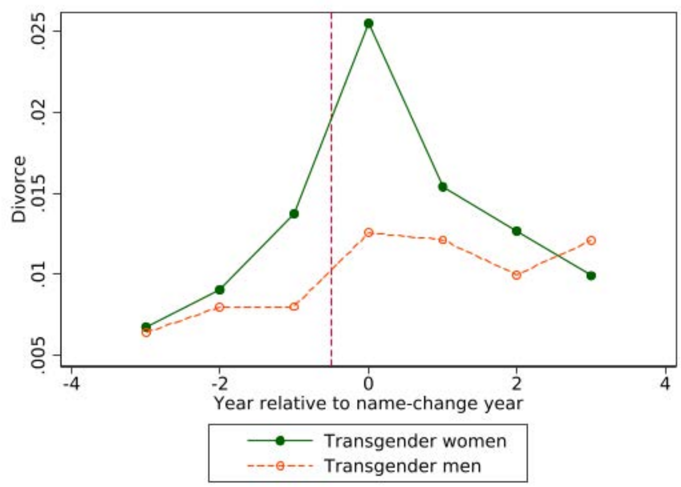

Notes: Author calculations from confidential IRS data. Each sample is reweighted to match the average distribution by age and name-change year. We restrict to individuals whose name-change year is 2019 or earlier.

F. Decomposition of transgender earnings gap in panel

In Appendix Table F1, we estimate the share of the total earnings gap from Figure 3 that is accounted for by the extensive margin. To do so, we define $y_{i\tau e}$ as "potential" earnings for person i in cohort τ at event time e – i.e., the earnings that that individual would have if they worked. Let $D_{i\tau e}$ denote whether a given individual works (i.e., has positive earnings) in the period. We observe $y_{i\tau e}$ only for those with positive earnings (i.e., $D_{i\tau e} = 1$).

The Poisson regression coefficient in Figure 3 for event time +4 recovers the difference-in-differences between $e = -4$ and $e = 4$ (which we denote as $\Delta\Delta$) of log $(E(y_{i\tau e} D_{i\tau e}))$. We define the intensive margin effect to be $\Delta\Delta \log\big(E(y_{i\tau e})\big)$ – i.e., the treatment effect if everyone were working in all periods – and the extensive margin to be what remains. The estimation challenge is that we do not observe y_{it} for individuals where $D_{i\tau e} = 0$, and thus we must "impute" it using the sample of those with $D_{i\tau e} = 1$. Specifically, among those with $D_{i\tau,4} = 1$, we use OLS to regress $y_{i\tau,4}$ on a quadratic in age, fixed effects for τ, and fixed effects for transgender status interacted with ten bins of observed earnings at event time -4 (i.e., $D_{i\tau,-4} \times y_{i\tau,-4}$), including a bin for those with zero earnings. We assign imputed earnings, $y_{i\tau,4}^{imp}$, to be equal to $y_{i\tau,4}$ for those with $D_{i\tau,4} = 1$ and $\hat{y}_{i\tau,4}$ for those with $D_{i\tau,4} = 0$. We repeat the same process in reverse to impute earnings for event time -4. Finally, we estimate the difference in differences in $\log\left(E\big(y_{i\tau e}^{imp}\big)\right)$ using the Poisson specification in Equation (1), restricted to event times -4 and 4; this represents an estimate for the intensive margin component of the total difference-in-differences estimate. We perform this entire procedure separately by sex assigned at birth.

Appendix Table F1 reports these results. Row 1 reports the total effect; rows 2 and 3 split this into the estimated intensive margin effect and the residual extensive margin effect. We estimate that approximately 44% of the penalty (-

0.078 out of -0.177) for transgender women is accounted for by the extensive margin (column 1). In column 2, we find that over 100% of the much smaller penalty (-0.055 out of -0.037) for transgender men is explained by the extensive margin. Put differently, this analysis suggests that transgender men and women face a relatively similar extensive margin penalty, while transgender women face a much larger intensive margin penalty.

In Appendix Figures F1-F6, we investigate further how individuals without earnings reported in Form 1040 may be supporting themselves. In particular, we plot stacked event studies where the dependent variable is a dummy that equals one when an individual is not working *and* meets some other condition – for example, having investment income. Appendix Figures F1-F3 study transgender women, while Appendix Figures F4-F6 study transgender men. In Appendix Figures F1 and F4, we focus on sources of government support that we can observe in the tax data: retirement-age Social Security income (OASI), Social Security Disability Income (SSDI), and Unemployment Insurance (UI). We see in Appendix Figures F1 that transgender women increase the prevalence of not working and receiving SSDI by nearly 1.5 percentage points relative to same-aged cisgender men; this represents a quantitatively modest share of the 9 percentage point overall extensive margin effect for transgender women. We see no meaningful effect in other government support nor non-labor market activities (Appendix Figures F2 and F5) and in non-labor, non-government sources of income (Appendix Figures F3 and F6). We note, however, that there are many components of the social safety net that we do not observe in tax data, including benefits through the Temporary Aid for Needy Families (TANF) and Supplemental Nutrition Assistance Program (SNAP).

Next, we decompose the remaining *intensive margin* changes into a component that can be "explained" by within-person changes in covariates and a residual "unexplained" component. To do so, we restrict to the set of individuals who have positive wages in both event times -4 and 4. In this sample, we can consider a variety of covariates – both those defined at the individual level, and those defined at the level of their highest-paying employer. In particular, the covariates we consider are: a dummy for the individual currently being enrolled in post-secondary education, experience and its square, a dummy variable for residing in a metropolitan area, the log density of the individual's zip code, and the Democratic vote share of the individual's county. We also consider the total number of the individual's employers, tenure at the employer, average log wages at the individual's firm, and average share of workers who are men at the individual's firm. The amount explained by covariates is given by $(\Delta\Delta X)\beta$, where $\Delta\Delta$ is the difference-in-differences operator (i.e., transgender minus cisgender earnings, post- minus pre-name change). Estimating $\Delta\Delta X$ is straightforward; to estimate β we run a Poisson regression of earnings on X, controlling for person fixed effects and year fixed effects, so that β is identified off of within-person changes. We do this regression separately in the cisgender panel and the transgender panel, so we have two estimates of β (and thus two estimates of $(\Delta\Delta X)\beta$).

We present the results from this intensive margin decomposition in Appendix Table F2; this table restricts attention to transgender women, as we do not estimate there to be an intensive margin penalty for transgender men. For each variable, we report (across the columns) the difference-in-differences estimate of the characteristic across cisgender and transgender samples before and after the name change (column 1), the association of each variable with cisgender earnings (column 2), the association of the variable with transgender earnings (column 3), the share of the intensive margin gap that is explained using the cisgender β (column 4), and the share of the intensive gap that is explained using the transgender β (column 5). Note the sign convention: if the change in covariates fully explained the intensive margin gap, columns 4 and 5 would sum to one.

The results in Appendix Table F2 do not return evidence that any single characteristic is driving the sizable intensive margin earnings difference between transgender women and cisgender men. There is some evidence that "job hopping" does contribute to the transgender earnings penalty: the combination of the number of employers and the job tenure effect explains 5 to 10 percent of the total effect, depending on which set of β estimates are used. Additionally, the experience quadratic explains 6 percent (using cisgender β) to 36 percent (using transgender β) of the total effect. Several other variables push in the "wrong" direction, however. For example, although average wages at one's firm is strongly positively associated with one's own wages, this does not explain the transgender earnings penalty in part because transgender individuals move to slightly higher paying firms on average. Overall, the main takeaway in Appendix Table F2 is that the amount of the earnings penalty between transgender women and cisgender men that can be explained by these covariates is small relative to the approximately 10 log point intensive margin earnings gap found in Appendix Table F1.

Appendix Table F1: Decomposition of Panel Earnings Gap into Extensive and Intensive Margins

	Transgender women	Transgender men
Total effect	-0.179	-0.037
Intensive margin	-0.101	0.027
Extensive margin	-0.078	-0.064

Author calculations from confidential IRS data. This table decomposes the total effect (i.e., the event time 4 coefficients from Figure 3) into the intensive margin effect and extensive margin effect, using the methodology described in Appendix F.

Appendix Table F2: Intensive Margin Decomposition of Earnings Gap Between Transgender Women and Cisgender Men

	(1)	(2)	(3)	(4)	(5)
	Difference in differences in covariate	Effect of covariate on earnings, cisgender sample	Effect of covariate on earnings, transgender sample	Amount explained (using cisgender β)	Amount explained (using transgender β)
Student	-0.045	-0.266	-0.296	-0.141	-0.157
Experience	-0.329	0.115	0.199	0.445	0.774
Experience squared	-8.950	-0.004	-0.004	-0.380	-0.417
Lives in a metro area	0.065	0.030	0.057	-0.023	-0.044
Log zip code density	0.345	-0.014	-0.023	0.055	0.094
Democrat vote share in zip code	0.036	0.082	0.412	-0.035	-0.175
Number of employers	0.407	0.020	0.046	-0.094	-0.223
Tenure	-1.176	0.010	0.023	0.142	0.313
Avg log wages at firm	0.013	0.279	0.346	-0.043	-0.054
Avg share male at firm	-0.034	-0.059	0.157	-0.024	0.063
Trans N	3,720	3,720	3,720	3,720	3,720

Author calculations from confidential IRS data. This table estimates the intensive margin earnings differences between transgender women before and after name-change (relative to cisgender men) that can be explained by changes in covariates, using the methodology described in Appendix F. Column 1 reports the difference-in-differences in each covariate, defining event time -4 as the pre-period and event time +4 as the post-period; the "treatment" group is transgender women while the "control" group is cisgender men. Column (2) reports the coefficients from a single Poisson regression of earnings on the listed covariates, controlling for person-by-cohort and post-by-cohort fixed effects (where cohort is defined in Section 4), estimated on the cisgender men in the control group. Column (3) reports coefficients from an analogous regression estimated on the transgender women in the treatment group. Columns (4) and (5) report the share of the total intensive margin effect explained by these changes in covariates – i.e., column (1) multiplied by column (2) or column (3), divided by the total intensive margin effect

Appendix Figure F1: Non-participation by transgender women: transfer income

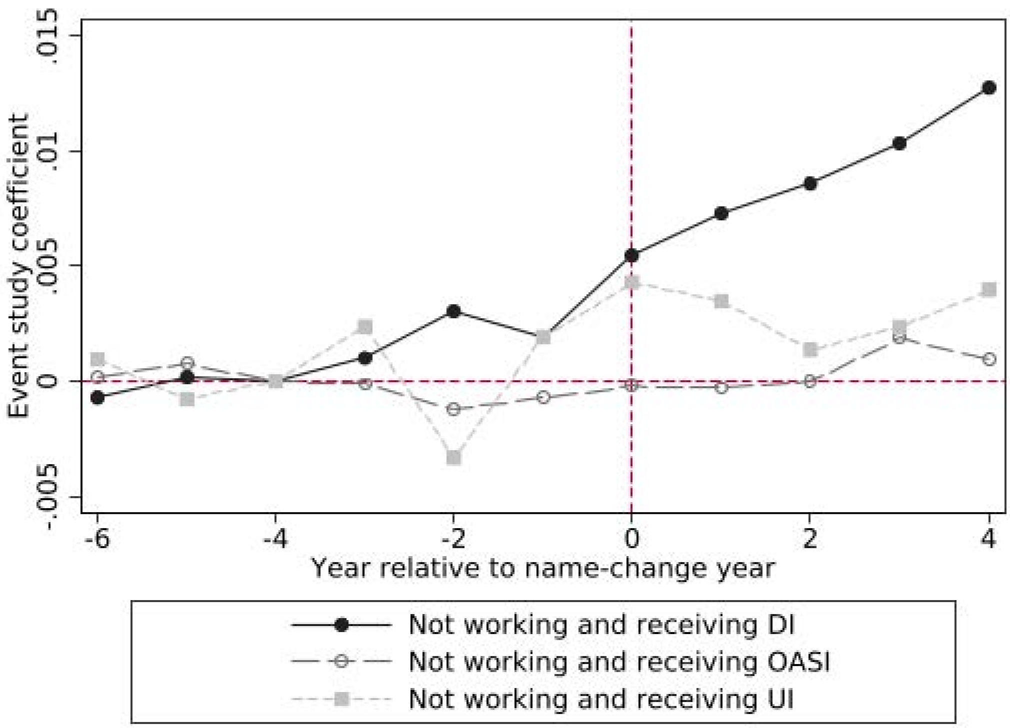

Notes: Author calculations from confidential IRS data. This figure plots stacked event study coefficients for transgender women (relative to cisgender men) for a set of three dependent variables. The dependent variables are dummies that equal one if and only if (a) an individual has zero earnings in a given year and (b) the individual satisfies some other condition. In this plot, the other conditions are (1) receiving Social Security Disability Insurance (DI) income, (2) receiving Old- Age & Survivor's Insurance (OASI) income, or (3) receiving unemployment insurance (UI) income. Event time -4 is omitted.

Appendix Figure F2: Non-participation by transgender women: family support and non-labor activities

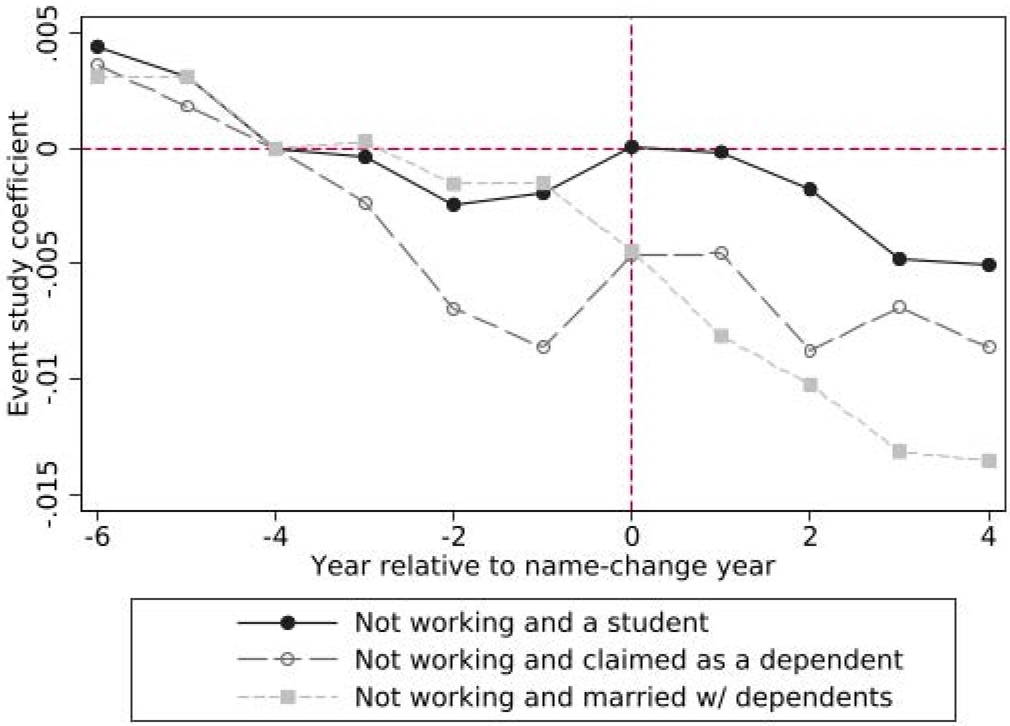

Notes: Author calculations from confidential IRS data. This figure plots stacked event study coefficients for transgender women (relative to cisgender men) for a set of three dependent variables. The dependent variables are dummies that equal one if and only if (a) an individual has zero earnings in a given year and (b) the individual satisfies some other condition. In this plot, the other conditions are (1) being a student, identified using receipt of Form 1098-T, (2) being claimed a dependent by someone else, or (3) being married with one or more dependents.

Appendix Figure F3: Non-participation by transgender women: other sources of non-transfer income

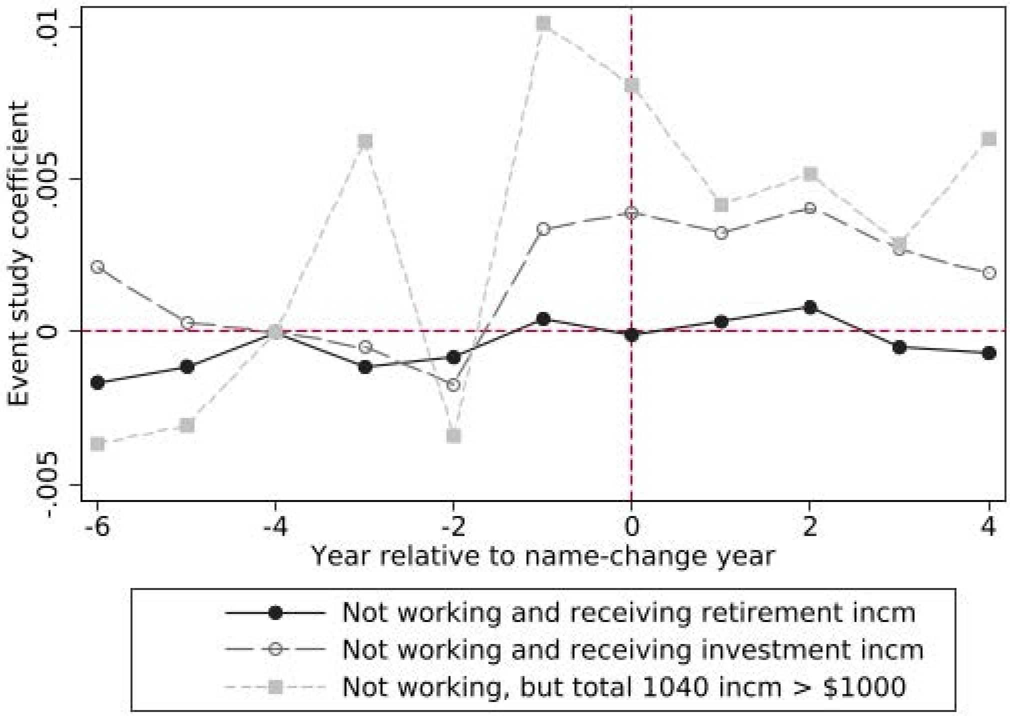

Notes: Author calculations from confidential IRS data. This figure plots stacked event study coefficients for transgender women (relative to cisgender men) for a set of three dependent variables. The dependent variables are dummies that equal one if and only if (a) an individual has zero earnings in a given year and (b) the individual satisfies some other condition. In this plot, the other conditions are (1) receiving at least $1000 of retirement income (defined as pension or IRA distributions) (2) receiving at least $100 of taxable interest or dividend income, or (3) having total income from Form 1040 of at least $1000. Event time -4 is omitted.

Appendix Figure F4: Non-participation by transgender men: transfer income

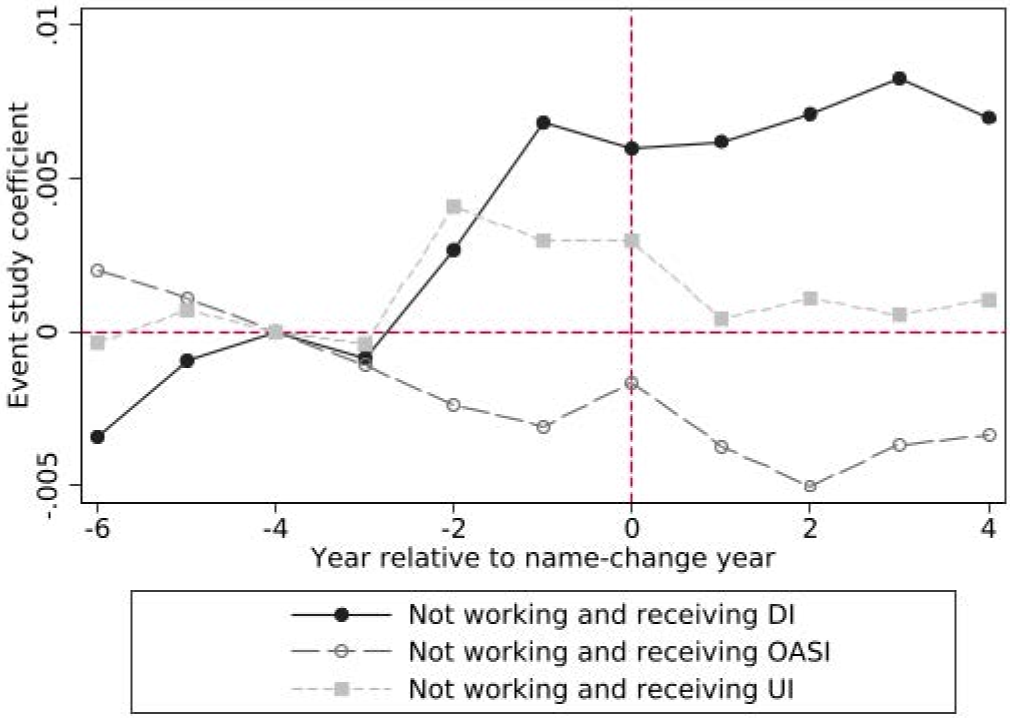

Notes: Author calculations from confidential IRS data. This figure plots stacked event study coefficients for transgender men (relative to cisgender women) for a set of three dependent variables. The dependent variables are dummies that equal one if and only if (a) an individual has zero earnings in a given year and (b) the individual satisfies some other condition. In this plot, the other conditions are (1) receiving Social Security Disability Insurance (DI) income, (2) receiving Old- Age & Survivor's Insurance (OASI) income, or (3) receiving unemployment insurance (UI) income. Event time -4 is omitted.

Appendix Figure F5: Non-participation by transgender men: family support and non-labor activities

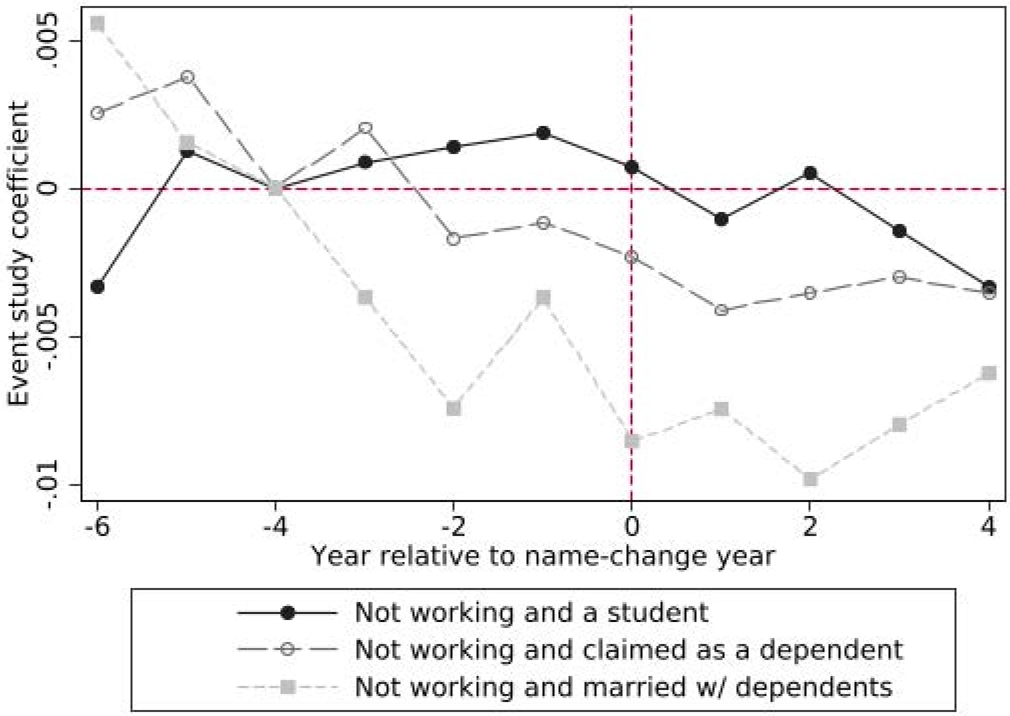

Notes: Author calculations from confidential IRS data. This figure plots stacked event study coefficients for transgender men (relative to cisgender women) for a set of three dependent variables. The dependent variables are dummies that equal one if and only if (a) an individual has zero earnings in a given year and (b) the individual satisfies some other condition. In this plot, the other conditions are (1) being a student, identified using receipt of Form 1098-T, (2) being claimed a dependent by someone else, or (3) being married with one or more dependents.

Appendix Figure F6: Non-participation by transgender men: other sources of non-transfer income

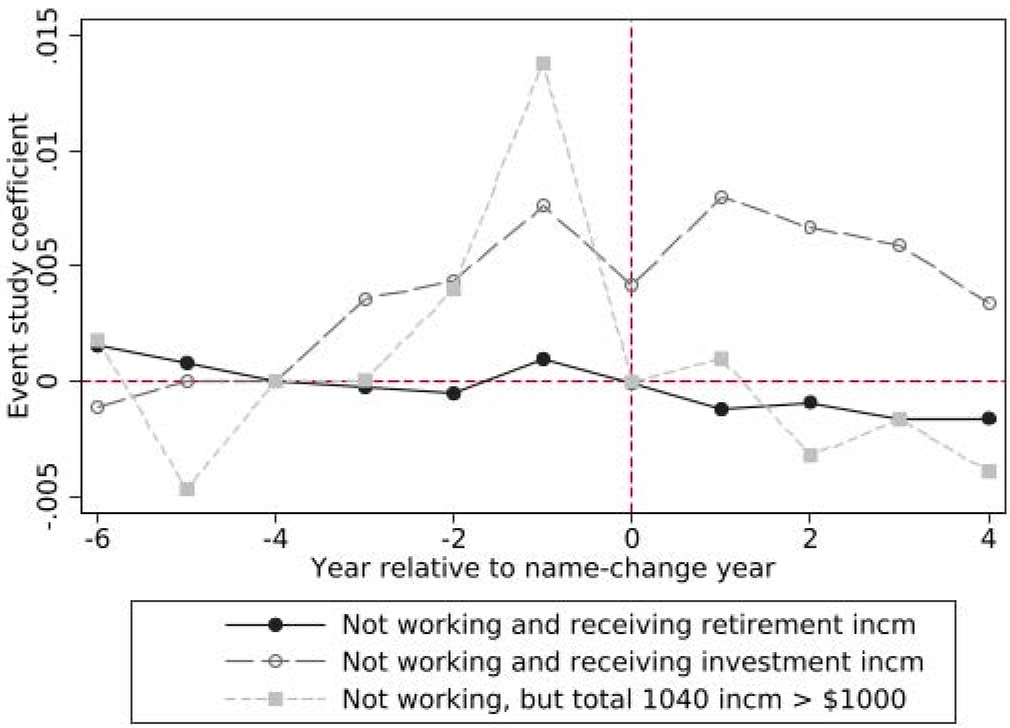

Notes: Author calculations from confidential IRS data. This figure plots stacked event study coefficients for transgender men (relative to cisgender women) for a set of three dependent variables. The dependent variables are dummies that equal one if and only if (a) an individual has zero earnings in a given year and (b) the individual satisfies some other condition. In this plot, the other conditions are (1) receiving at least $1000 of retirement income (defined as pension or IRA distributions) (2) receiving at least $100 of taxable interest or dividend income, or (3) having total income from Form 1040 of at least $1000. Event time -4 is omitted.

G. Construction of industry and occupation variables

Regarding industry, we use the employer-reported NAICS codes on business tax returns (including their payroll tax returns, meaning that government and non-profit employers are also included). In some cases, the IRS fills in missing NAICS codes when businesses are audited. For each individual, we attach the NAICS code associated with the business that paid them the most wages in a given year. We treat NAICS codes as missing when the implied two-digit NAICS is invalid. When the NAICS code is missing for an individual in a given year – e.g., because of non-employment, or because the highest-paying employer did not have a valid NAICS code – we impute forward from the most recent non-missing year. Among transgender individuals aged 20 and above, 98% have a valid industry (after the aforementioned imputation) in 2022.

Regarding occupation, we use the free-text occupation field on Form 1040, which we observe for those who file electronically. We use the crosswalk from that free-text field to SOC codes developed by Bruce Sacerdote and sent to us via personal correspondence; this crosswalk accounts for common misspellings. When the SOC code is missing – either because of non-filing, paper filing, or because the free text occupation field is not informative – we impute forward from the most recent non-missing year. Among transgender individuals aged 20 and above, 81% have non-missing occupation (after the aforementioned imputation) in 2022.

Appendix Table G1: Industry

	(1)	(2)	(3)	(4)
Industry	Transgender men	Cisgender men	Transgender women	Cisgender women
Agriculture, Forestry, Fishing and Hunting	0.006	0.014	0.004	0.006
Mining, Quarrying, and Oil and Gas Extraction	0.001	0.005	---	0.001
Utilities	0.001	0.003	---	0.001
Construction	0.016	0.080	0.011	0.014
Manufacturing	0.049	0.102	0.070	0.051
Wholesale Trade	0.018	0.033	0.021	0.020
Retail Trade	0.177	0.124	0.151	0.105
Transportation and Warehousing	0.025	0.042	0.025	0.018
Information	0.023	0.020	0.057	0.018
Finance and Insurance	0.024	0.026	0.030	0.040
Real Estate and Rental and Leasing	0.015	0.018	0.012	0.019
Professional, Scientific, and Technical Services	0.092	0.086	0.148	0.096
Management of Companies and Enterprises	0.023	0.026	0.027	0.024
Admin. & Support & Waste Management Services	0.082	0.100	0.083	0.086
Educational Services	0.056	0.019	0.038	0.039
Health Care and Social Assistance	0.094	0.037	0.067	0.163
Arts, Entertainment, and Recreation	0.025	0.020	0.021	0.015
Accommodation and Food Services	0.085	0.092	0.066	0.079
Other Services (except Public Administration)	0.050	0.038	0.039	0.050
Public Administration	0.139	0.115	0.127	0.157

Author calculations of sample means in 2022 from confidential IRS data. The sample of cisgender men (women) is a random sample stratified by cohort of those who filed a 2022 tax return, appeared as a dependent on a 2022 tax return, or had an information return in 2022. The sample of cisgender men (women) has been re-weighted to match the year and cohort distribution of transgender men (women). Industries are defined by two digit NAICS codes. We restrict attention to individuals with a valid two-digit industry and positive wages in 2022. Cells marked with "---" are suppressed for disclosure avoidance. See Appendix G for details on the construction of the industry variable.

Appendix Table G2: Occupation

	(1)	(2)	(3)	(4)
Occupation	Transgender men	Cisgender men	Transgender women	Cisgender women
Management	0.068	0.075	0.067	0.082
Business and Financial Operations	0.042	0.049	0.047	0.063
Computer and Mathematical	0.039	0.043	0.172	0.018
Of which: Software Developers, Applications	*0.013*	*0.019*	*0.084*	*0.005*
Architecture and Engineering	0.014	0.033	0.043	0.010
Life, Physical, and Social Science	0.015	0.009	0.016	0.009
Community and Social Service	0.043	0.008	0.022	0.027
Legal	0.009	0.005	0.009	0.011
Education, Training, and Library	0.057	0.021	0.045	0.082
Arts, Design, Entertainment, Sports, and Media	0.042	0.018	0.045	0.023
Healthcare Practitioners and Technical	0.054	0.025	0.033	0.098
Healthcare Support	0.023	0.007	0.016	0.051
Protective Service	0.028	0.033	0.017	0.011
Food Preparation and Serving Related	0.106	0.075	0.079	0.066
Building and Grounds Cleaning and Maintenance	0.015	0.023	0.012	0.021
Personal Care and Service	0.042	0.016	0.031	0.065
Sales and Related	0.130	0.099	0.114	0.109
Office and Administrative Support	0.134	0.087	0.115	0.162
Farming, Fishing, and Forestry	0.002	0.005	0.001	0.001
Construction and Extraction	0.015	0.073	0.010	0.004
Installation, Maintenance, and Repair	0.018	0.059	0.017	0.004
Production	0.033	0.068	0.025	0.028
Transportation and Material Moving	0.061	0.139	0.054	0.045
Military	0.010	0.030	0.010	0.006

Author calculations of sample means in 2022 from confidential IRS data. The sample of cisgender men (women) is a random sample stratified by cohort of those who filed a 2022 tax return, appeared as a dependent on a 2022 tax return, or had an information return in 2022. The sample of cisgender men (women) has been re-weighted to match the year and cohort distribution of transgender men (women). Occupations are defined by two digit 2010 SOC codes, except that "Software Developers, Applications" is a six digit 2010 SOC code; this row is also included in the "Computer and Mathematical" row. We restrict attention to individuals with a valid SOC codes. See Appendix G for details on the construction of the occupation variable.

H. The role of different control variables in the siblings comparisons

In this Appendix, we provide further detail into the effect of various control variables in affecting the measured transgender earnings gap using the siblings comparison from Section IV.B. In Appendix Table H1, we estimate the Poisson regressions described in that section using various sets of controls. Columns 1 and 2 repeat columns 1 and 2 of Table 2 – using no controls, and education and potential experience controls, respectively.

Column 3 of Appendix Table H1 adds marriage and family controls (a dummy for being married, a dummy for having a dependent 5 or younger, and a dummy for having a dependent 6 or older); these controls make the transgender woman penalty (row 1) much less negative and has a slightly negative effect on the transgender man penalty (row 2), changing it from positive to negative. Column 4 adds geography controls (the log of the zip code's population density, and a dummy for being a central county in one of the top 30 CBSAs), making the transgender penalty more negative for both genders. Column 5 adds industry and occupation fixed effects, arriving at column 3 of Table 2, which again makes the penalty more negative. These two columns reflect the fact that transgender individuals tend to sort into urban areas (paying higher wages) and higher-earning industries and occupations.

Finally, Column 6 of Appendix Table H1 adds several variables relating to labor market outcomes: the number of years of tenure at each individual's highest-paying job[43], a quadratic in actual (as opposed to potential) experience, and the count of different employers from 1999 onward.[44] Including these variables as controls makes the transgender women penalty slightly more negative and transgender women penalty modestly less negative.

[43] We define tenure to be zero for those not earning any wages.

[44] The latter is motivated by evidence that life cycle wage gains are often driven by job-to-job transitions, e.g. Haltiwanger, Hyatt, and McEntarfer (2018).

We note that adding controls to these regressions can have multiple effects. In the simplest case, a certain covariate may reflect a labor-market relevant, exogenous characteristic that tends to differ between transgender individuals and their cisgender siblings; including such a control can help unconfound the regression. Broadly following the gender gap literature, we think of education and potential experience controls as mostly fitting this mold, though we acknowledge that education choices may be influenced by labor market opportunities.

By contrast, some controls are "bad controls" in the sense of Cinelli, Forney, and Pearl (2022). This includes cases where the control is a *mediator* – i.e., a mechanism through which labor discrimination (or some other factor causally driving the transgender earnings gap upward or downward) operates. This likely applies to some extent to the geography controls, industry and occupation fixed effects, and certainly to the labor market variables. Thus, including these controls can shed some light on which mechanisms matter quantitatively; the fact that geography and industry and occupation fixed effects drive the penalty more negative implies that sorting along these dimensions cannot explain the gap. By contrast, some of the labor market variables (including tenure and actual experience) can explain a modest share of the gap for transgender women, but not transgender men.

Finally, in our context, including marriage and childrearing as controls may have somewhat unintuitive effects. Marriage tends to be positively correlated with outcomes (especially for cisgender men), reflecting a selection effect and possibly a causal effect. At the same time, transgender individuals are much less likely to be married, perhaps because of differing attitudes toward marriage as an institution or differing preferences and costs regarding childrearing. This can cause marriage and childrearing to act as a "collider" in the regression, biasing the regression coefficient of interest upward (toward zero). For example, suppose there is a latent variable U that is positively correlated with both earnings and marriage. If marriage is "costlier" for transgender individuals, then married transgender individuals will tend to have higher U draws than married cisgender individuals – and the same is true for unmarried transgender individuals relative to unmarried cisgender individuals. This creates the possibility for collider bias. Nevertheless, as seen in Column 3 of Appendix Table H1, this bias is not so severe as to eliminate the entirety of the transgender gap.

Appendix Table H1: Siblings Fixed Effects Estimates of Transgender Earnings Gaps

	No controls (Sibling pair fixed effects only)	+ Education and potential experience controls	+ Marriage and family controls	+ Geography controls	+ Industry and occupation fixed effects	+ Labor market variables
			Transgender women and cisgender brother sibling pairs			
Transgender	-0.173 (0.023)	-0.192 (0.022)	-0.099 (0.023)	-0.144 (0.024)	-0.217 (0.021)	-0.159 (0.023)
			Transgender men and cisgender sister sibling pairs			
Transgender	0.017 (0.018)	0.019 (0.017)	-0.016 (0.019)	-0.025 (0.019)	-0.038 (0.018)	-0.057 (0.017)
			All transgender and cisgender sibling pairs			
Transgender	-0.079 (0.011)	-0.083 (0.011)	-0.044 (0.011)	-0.073 (0.011)	-0.130 (0.010)	-0.118 (0.012)
Male	0.095 (0.015)	0.113 (0.015)	0.104 (0.014)	0.119 (0.015)	0.133 (0.013)	0.079 (0.015)
Male at birth	0.202 (0.015)	0.188 (0.015)	0.205 (0.015)	0.192 (0.015)	0.062 (0.015)	0.077 (0.017)
Trans N	20,870	20,870	20,870	20,870	20,870	20,870

Author calculations from confidential IRS data. This table reports estimates using a sample of transgender individuals linked to their siblings, as discussed in Section IV.B. We restrict to transgender individuals and siblings within the 1981-1999 birth cohorts. “Education and potential experience controls” includes a linear term for number of years of education, a dummy for currently being a college student, a dummy for currently being enrolled more than half time, and a quadratic in potential experience. “Marriage/family controls” include a dummy for being married, a dummy for having a youngest dependent under age 6, and a dummy for having a youngest dependent age 6 and above. “Geography controls” include log population density of the zip code, and a dummy for being in a “central county” in one of the top 30 MSAs ranked by population. “Industry/occupation fixed effects” are fixed effects for two-digit NAICS code of the highest-wage employer (including missing) and fixed effects for two-digit occupation (including missing). “Labor market variables” are the number of years of tenure at each individual’s highest-paying job (defined as zero for those without wage income), a quadratic in actual (as opposed to potential) experience, and the count of different employers from 1999 onward. Standard errors are clustered by unique transgender individuals. For disclosure-avoidance purposes, we round sample counts to the nearest multiple of 10.

I. Additional Tables and Figures

Appendix Figure I1: Log mean earnings, by group and event time

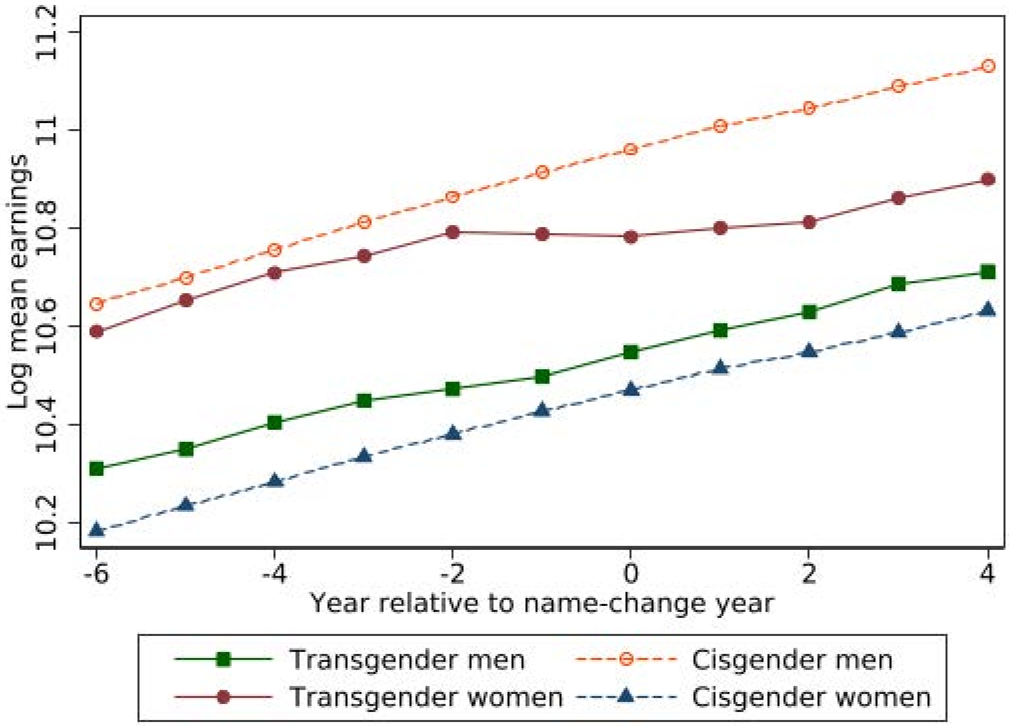

Author calculations from confidential IRS data. This figure plots the log of mean earnings by event time for each of the four groups defined by transgender/cisgender men/women, using the stacked samples included in the stacked event study regression in Figure 3. Each sample (transgender women or men and cisgender men or women) is reweighted to match the age distribution of transgender individuals (aggregated between transgender men and transgender women). We restrict to transgender individuals whose age at name-change year is 28 or greater, and whose name-change year is 2018 or earlier. See Section IV.A for further implementation details.

Appendix Table I1: Panel Estimates of Transgender Earnings Gaps

Event time	Transgender men coefficient	Transgender women coefficient
-6	0.006 (0.012)	-0.009 (0.019)
-5	-0.008 (0.009)	0.001 (0.011)
-4	--	--
-3	-0.005 (0.010)	-0.020 (0.014)
-2	-0.029 (0.014)	-0.022 (0.023)
-1	-0.049 (0.015)	-0.077 (0.025)
0	-0.042 (0.016)	-0.129 (0.023)
1	-0.041 (0.016)	-0.157 (0.026)
2	-0.037 (0.020)	-0.181 (0.025)
3	-0.019 (0.031)	-0.176 (0.026)
4	-0.037 (0.019)	-0.179 (0.026)
Total sample size	17,061,380	17,403,120
Clusters	217,800	221,780
Trans N	3,720	6,440

Author calculations from confidential IRS data. This table reports estimates from three Poisson stacked event study regressions described in Section IV.A of the text. Each column corresponds to a separate regression. In the first column, transgender men are treated and cisgender women are control. In the second column, transgender women are treated and cisgender men are control. We restrict to transgender individuals whose age at name-change year is 28 or greater, and whose name-change year is 2018 or earlier. See Section IV.A for further implementation details. Standard errors are clustered by unique individual. For disclosure-avoidance purposes, we round sample counts to the nearest multiple of 10.

Appendix Table I2: Covariate Differences Between Transgender Individuals and Their Siblings

	Transgender individuals	Cisgender siblings
Number of years of college	5.01	4.58
Currently a student	0.170	0.139
Potential experience	7.37	7.83
Married	0.203	0.322
Has a dependent	0.078	0.276
Log density of zip code	7.73	7.28
Lives in central city of top 30 MSA	0.525	0.455
Weighted sample size	21,350	21,350

Author calculations of sample means from confidential IRS data in 2022. Both columns restrict to those individuals used in the sibling analysis described in Section IV.B. See Section IV.B for additional sample details. For disclosure-avoidance purposes, we round sample counts to the nearest multiple of 10.

Appendix Table I3: Covariate Differences Between Transgender Individuals and Their Coworkers

	Transgender individuals	Cisgender coworkers
Number of years of college	5.197	4.462
Currently a student	0.150	0.123
Potential experience	7.610	8.418
Married	0.231	0.317
Has a dependent	0.086	0.307
Log density of zip code	7.780	7.503
Lives in central city of top 30 MSA	0.527	0.499
Weighted sample size	13,230	13,230

Author calculations of sample means from confidential IRS data in 2022. Both columns restrict to those individuals used in the coworker analysis described in Section IV.C. See Section IV.C for additional sample details. For disclosure-avoidance purposes, we round sample counts to the nearest multiple of 10.