The male breadwinner norm in Brazil: a bunching approach

Kauany Souza
University of São Paulo
kalsouza@usp.br

Fabiana Rocha
University of São Paulo
frocha@usp.br

Pedro Forquesato
University of São Paulo
pforquesato@usp.br

We investigate the impact of the male breadwinner norm – a cultural norm that establishes that "a man should earn more than his wife" – on married females earning outcomes. Our work aims to verify if this identity norm shapes behavior, and we start by modelling it as an implicit tax on women's earnings. We estimate the labor supply response of a wife who is likely to just out-earn her husband and reacts by adjusting her earnings so as to not violate the gender norm. Using bunching techniques and a range of labor supply elasticities, we find that the breadwinner norm generates a sizable additional tax rate on female labor supply. Our preferred estimate suggests the social norm add a 34 percentage points implicit marginal tax on earnings for those women who out-earn their husbands with similar earnings. Additional extensions suggest that religiosity and age play a role on the implicit tax related to this gender norm.

1 Introduction

Gender norms – understood as a set of beliefs and social expectations characterizing masculinity and femininity (RISMAN, 1999) – are widely discussed but remain understudied in quantitative terms. They particularly affect conjugal relationships, as couples whose behaviors and arrangements deviate from prevailing norms become more likely to experience social sanctions, relationship frictions and stigma (KILLEWALD, 2016).

The aim of this article is to analyse the male breadwinner norm, a gender norm that is both very common across cultures and increasingly debated, which states that a man should always earn more than his female partner. Our analysis will focus on the tax-like effects that adjustments to the male breadwinner norm have on female earnings, and it's empirical design relies primarily on bunching techniques (Kleven, 2016) over a constructed record of brazilian couples surveyed by the IBGE's *Censo Demográfico*.

In Brazil, the World Values Survey (2020) found that more than 35.5% of respondents agree that "it's a problem if a woman have more income than her husband"; 14.5% neither agrees or disagrees. This openly-held belief ought to affect labor market incentives similarly to prices and technology – gender norms are, as described by the gender studies literature, powerful social constructs that deeply shape human behavior (CONNELL, 2010).

Submitted to:
Anpec 2022

There has been a recent surge of interest on the male breadwinner norm in economics, which mostly stems from observations of an empirical occurrence seem in several countries: the population distribution of relative spousal earnings tends to show a discontinuity around the point where women start out-earning their husbands. This phenomenon also occurs in Brazil, as can be seem in Figure 1.

So as to start our analysis, we'll model the male breadwinner norm as a social disutility which penalizes every *real* of women's earnings when they out-earn their husbands. This has direct parallels with bunching analysis of notched tax rate schedules. Under reasonable assumptions about preferences, the male breadwinner norm generates a convex kink in women's allocation set which both explains bunching behavior around the 0.5 relative income threshold and allows us to non-parametrically estimate it's size and cost. Afterwards, we'll use earnings elasticity estimates to calibrate our results and find the implicit tax ascribed to this breadwinner norm.

Figure 1: Distribution of wife's earning shares – *Censo Demográfico 2010*

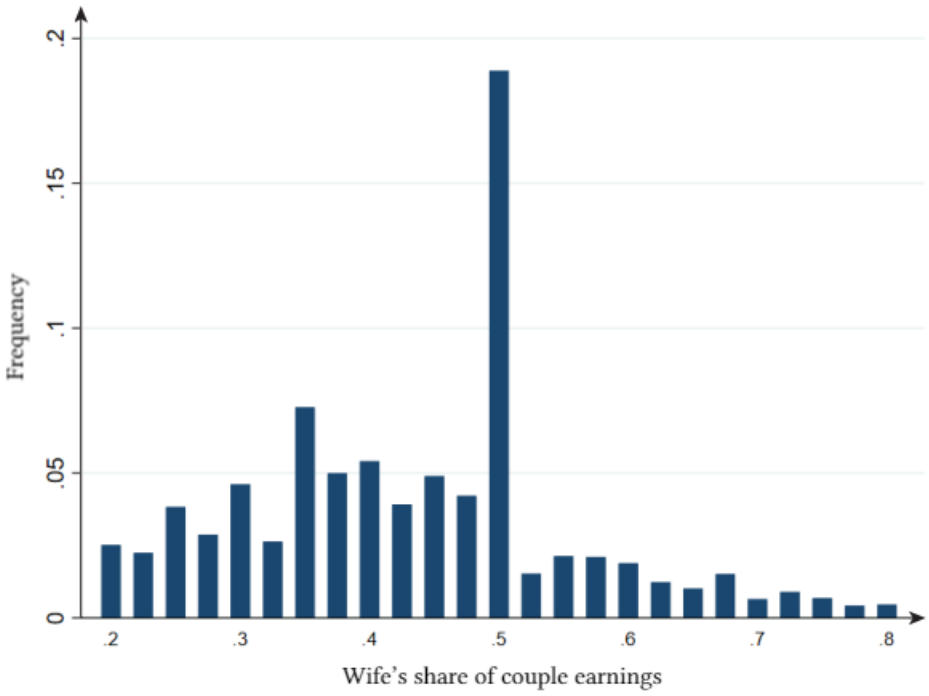

Our results are shown to be consistent with the presence of a male breadwinner norm in Brazil. We find a sharp bunching in the relative earnings distribution at the 0.5 threshold that cannot be explained by alternative sources of discontinuity. For our baseline specification, we find a money-metric tax rate on female earnings that range from 20% to 43% throughout different census decades. Religious affiliation and age also play a role: women appear to pay up to 12% more implicit gender norm tax over their earnings if they're religious compared to their non-religious counterparts, as well as older women may pay over 11% in those taxes over the younger married cohort. The estimates are highly robust to the choice of different parameters, and, given their size, have high economic relevance.

By considering gender norms in an economic model of behavior, and by estimating the monetary costs associated with them, we aim to make a contribution to two strands of literature. First, our work relates to the broad literature of economics and identity (AKERLOF; KRANTON, 2000), expanded in terms of behavioral responses to couples' earning inequalities by Bertrand et al. (2015). Second, our work contributes to the applied bunching methodology literature started by Saez (2010) in the context of taxation and that is now widespread to a wide range of non-tax applications as discussed by Kleven (2016). To our knowledge, this is the first bunching analysis of a social norm using Brazilian data to date.

The article is organized in four sections, besides this introduction. Section 2 initiates our literature review. Section 3 introduces our theoretical model of the breadwinner norm as a tax. Section 4 brings the empirical discussion and discusses our results for the Brazilian data. Finally, section 5 contains our concluding remarks.

2 Literature Review

Gender identity relates to a number of research areas in economics and sociology. This section surveys the emerging literature on the behavioral effects of gender norms in general and the male-breadwinner norm specifically. In economics, recent research have tried to relate the change in wife's relative income to a few specific outcomes in relationship dynamics, like increases in divorce rates, housework contribution, or decreases in female labor market participation, which all aggregate to indirect impacts on earnings themselves.

A starting point is the seminal work by Akerlof and Kranton (2000), which laid the foundation for the study and modelling of norms. The authors introduce the concept of socially determined identities into individuals' utility functions. In their modelling, a person's belonging into identity groups circumscribe a clear conception of how members of these groups should behave, and any deviation from those behaviors generates costs from social pressure to conform. In the case of a male breadwinner norm, gender identities' characterization of a breadwinner husband and homemaking wife would lead to a decline in the couple's utility as the wife marginally out-earns the husband.

In an alternative view, some authors choose to highlight the role of information or social learning (FOGLI; VELDKAMP, 2011; CAVAPOZZI; FRANCESCONI; NICOLETTI, 2021) on the adherence of the male breadwinner norm. As some women may be uncertain about the effects of changing relative earnings on their family's well-being, the quality of their relationships or their work-life balance, they may therefore look to other same-sex adults and to peers for valuable information. Both channels – social pressure or transmission – demonstrate how local and context-dependent the adoption of the male-breadwinner norm works.

Bertrand et al. (2015) provides prominent evidence on the existence of the male-breadwinner norm in the United States. The authors argue that the sharp cliff around the 0.5 mass point of the wife's relative earnings distribution, shown as a discontinuity through a McCrary (2008) test, share distinctive features with a distribution under the norm. They also show that, with appropriate controls, marriage

rates decline when the predicted probability that a woman earns more than the man increases. In an individual level, a higher probability that the wife's potential income exceeds her husband's actual income makes the wife less likely to participate in the labor force – and if she does work, the gap between her realized and potential incomes is higher. These patterns, the authors argue, suggest women distort their labor supply so as to avoid a break of the male breadwinner norm.

Multiple studies echo Bertrand et al.'s findings. Cooke (2006) finds that among couples in the US where the wife earns more than the husband, the likelihood of divorce is lower if the wife engages in "compensatory behavior", in which she does more household chores. Greenstein (2000) argues that both men and women adjust their domestic production to neutralize the cost of deviating from gender roles. Bianchi et al. (2000) found a negative linear relationship between a wife's relative income and the amount of time spent on housework in the US, while Bittman et al. (2003) found a U-shaped relationship in Australia — after the wife earns more than her husband, greater increases in her relative income translate into greater inequalities in housework.

Codazzi, Pero and Sant'Anna (2018) found evidence in Brazil that the 0.5 drop in the wife's share of household earnings is associated with wives distorting their labor supply in order not to violate the male breadwinner norm. The authors estimate the impact of higher probabilities of a woman out-earning her husband on many outcome variables, such as wife's labor participation, working hours and informal work. Their results pointed to significant and sharp positive effects. Similar estimates were also found in former West Germany (SPRENGHOLZ; WIEBER; HOLST, 2020), where norm-conforming labor distortions were salient before the unification and dwindled afterwards, and China (DONGCHENG; FANBO; ZIXUN, 2021), where women with higher potential income than their husbands' actual income accrued lower earnings and working hours (although no effects were found in their labor market participation).

Some studies have recently disputed Bertrand et al. (2015) and similar works by questioning the use of wife's relative earnings distribution as an identification strategy for the male breadwinner norm. Those objections fall broadly under the category of assortative matching considerations (BINDER; LAM, 2018), co-working and self-employed coupling (ZINOVYEVA; TVERDOSTUP, 2021) and misreporting (SLOTWINSK; ROTH, 2021). We'll catalog and address them in detail on the empirical section 4.

2.1 Bunching and labor supply

The presence of discontinuities in empirical distributions around the point at which agents are subject to considerable changes in utility as a result of small behavioral adjustments suggest that these patterns may reveal important information about the responsiveness of such agents. This is especially relevant when the changes in utility are related to a discontinuous function of the endogenous variable under analysis. Bunching techniques (Saez, 2010; Kleven and Waseem, 2013) were created to work under these circumstances – as agents may alter their choices to ensure that they remain on the desired side of a cutoff, we can use such behavior to non-parametrically uncover structural parameters related to

their responses.

Specifically, the bunching estimation aims to construct a measure of the excess mass of individuals at the point of incentive discontinuity – called a kink, for marginal increases in disutility, or a notch, for average increases – by locally comparing the density mass of individuals at the bunching point with the mass at the same point without the discontinuity, which is inferred from a counterfactual distribution constructed from the data. Using such information, along with the amount of change in incentives - usually in the form of changing tax rates -, the method can uncover local elasticity estimates.

Although the early bunching literature was developed exclusively in the context of taxation (Kleven, 2016), particularly piecewise linear income tax schedules, several non-tax settings have been published using the technique. Those include energy performance certificates (COLLINS; CURTIS, 2018), education (DIAMOND; PERSSON, 2016), labor regulation (GOURIO; ROYS, 2014), Maastricht treaty's fiscal criterion (CASELLI; WINGENDER, 2018), cellular service prices (GRUBB; OSBORNE, 2015), amongst others.

So as to add to this list of applications, we'll analyze the notched incentives associated with the male breadwinner norm, and, unconventionally, use our bunching estimates to find an implicit gender norm tax on female earnings associated with it. We'll achieve this by imputing the estimates over a range of labor supply elasticities on par with those verified in the literature for women and/or Brazilians.

There's a large variation of women's labor supply elasticity estimates across studies. In a literature survey of labor supply and taxation, Grubb e Osborne (2011) shows that such estimates are usually high for women in different types of models that include participation choices, however they can sometimes be small when those are not incorporated. Mattos and Terra (2016) explored a limited tax reform implemented in the late 90s in Brazil and found elasticity estimates for all genders with relative high ranges when including cash and in-kind transfers. We'll aim to encompass those elasticity estimates further in the empirical section.

3 Theoretical framework

A simple household model will be initially used to highlight how a male breadwinner norm incite patterns over a couple's relative income population distribution. In this section I'll abstract away from a few details that are relevant to the empirical setting for the sake of exposition.

3.1 Model

There are one type of decision-making agent, wives, that can derive utility from the consumption of a numeraire market good. Each wife has a unique exogenously given different-sex husband, and there are $w = 1, 2, 3...\mathbb{W}$ wives. Amongst them, there is heterogeneity in ability n captured by a density function $f(n)$. In particular, I assume that wives preferences are additively separable and can be described by a direct quasi-linear utility function:

$$u_w(z_w) = z_w - v(z_w; n, e) - T(z_w) \cdot 1\left[\frac{z_w}{z_m} > 1\right] \tag{1}$$

Where z_w and z_m are, respectively, the gross incomes of wife and husband, $v(z_w; n, e)$ represents the wife's effort disutility associated with acquiring z_w such as $v'(z_w; n, e) > 0$, $v''(z_w; n, e) \geq 0$, and $T(z_w) \cdot 1\left[\frac{z_w}{z_m} > 1\right]$ is a social disutility given by a break of the gender norm that a man should make more than his female partner, which works as a single marginal rate of discount on wife's gross income $T(z_w) = \tau z_w$. This introduces the notch, as the utility function is maximized over her choice of income. Following Kleven and Waseem (2013), we'll specify an isoelastic effort function, so as to rule out income effects for τ changes:

$$v(z_w; n, e) = \frac{n}{1 + \frac{1}{e}} \left(\frac{z_w}{n}\right)^{1+\frac{1}{e}} \tag{2}$$

Let $\theta = \frac{z_w}{z_m + z_w}$ denote the wife's share of couple earnings. Suppose a gender norm-free world where $\tau = 0$ (i.e. there's no social disutility related to who's breadwinner in a relationship).

Figure 2: Wife's income optimization diagram

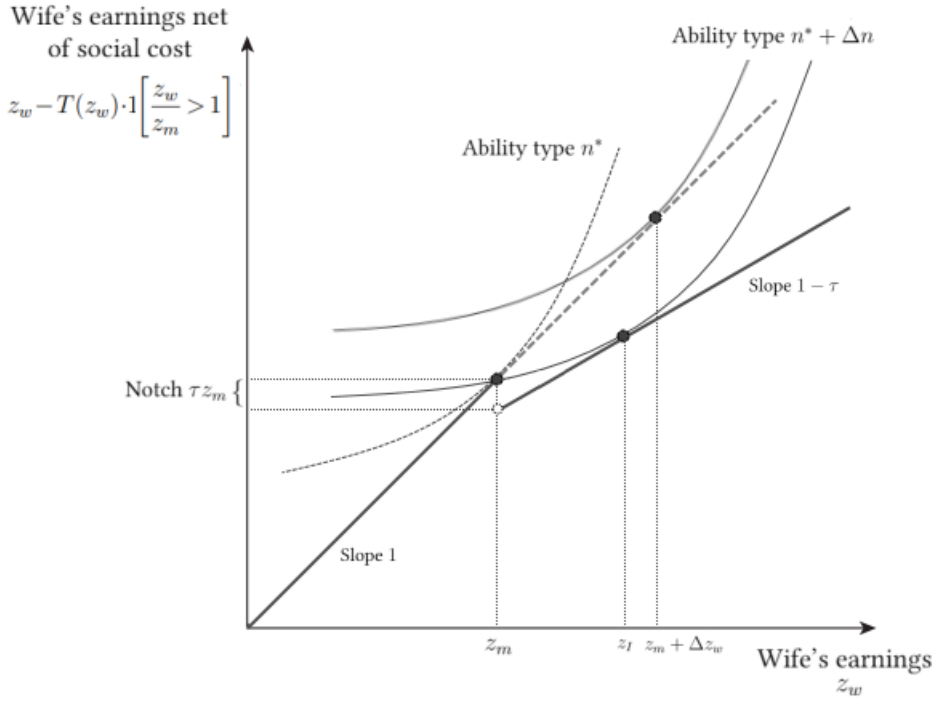

If we assume independence, smooth distribution of ability and of husband's incomes, wives' shares of their respective relationship combined earnings are also smoothly distributed according to a density function $h_0(\theta)$. If then the male breadwinner gender norm is introduced, so that for $z_w > z_m$ the unit

cost of additional earnings is τ, a change will take place at the relative earnings distribution, as some women will modify their gross earnings to adapt to the new norm.

Denote the density function of wives' relative income distribution after the gender norm is introduced by $h(\theta)$. For those women who used to make less than their husbands, first order effects of the norm will not affect them, so that $h(\theta) = h_0(\theta)$. Inversely, those who used to make more than their husbands will reduce their earnings given the new imposed social cost.

Figure 3: Relative income density distribution diagram

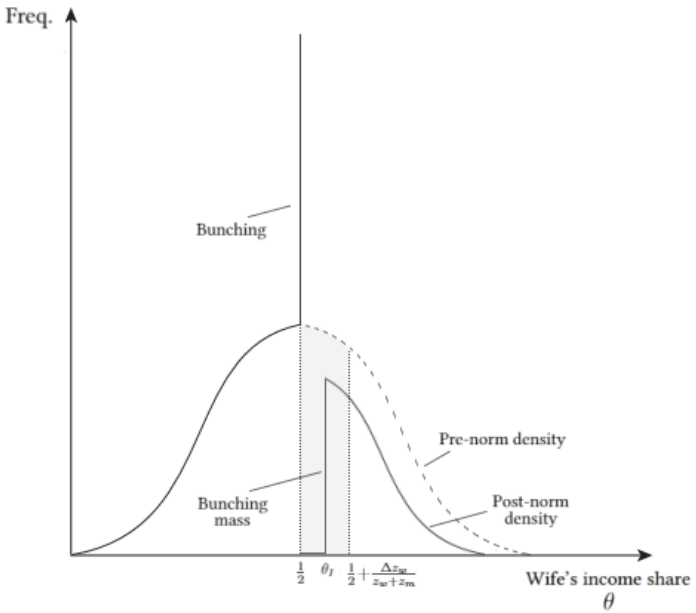

For a subset of women whose abilities $n \in [n^*, n^* + \Delta n]$ established their income before the gender norm into a range $z_w \in [z_m, z_m + \Delta z_w]$, the notched character of the new constraint will make them adjust to their husband's income z_m, and a spike of $\theta = 0.5$ occurs. There is a hole in the relative earnings density distribution right after the notch, as no wife is willing to earn between $[z_m, z_I]$. This can be properly visualized in Figures 1 and 2. Importantly, the wife originally located at the $z_m + \Delta z_w$ earnings point will be the marginal buncher wife, as she will be indifferent between locating at the notch point or at her best interior point z_I after the notch. The number of wives who make the adjustment to the notch will be numerically equivalent to the bunching mass in the relative earnings distribution:

$$B = \int_{\frac{1}{2}}^{\frac{1}{2}+\frac{\Delta z_w}{z_w+z_m}} h_0(\theta)\, d\theta \approx h_0\left(\frac{1}{2}\right)\frac{\Delta z_w}{z_m} \tag{3}$$

Importantly, as in Kleven (2016), we can relate the relative earnings response of the marginal wife buncher to the compensated elasticity e using her utility indifference at the husbands' earnings point $u_w^* = u_w(z_m)$ and at it's post-notch interior point $u_w^I = u_w(z^I)$. Given that $z^I = (1-\tau)^e(n+\Delta n)$ from first order conditions, we have:

$$\underbrace{\left(\frac{1}{1+e}\right)(n+\Delta n)(1-\tau)^{1+e}}_{u_w^I} = \underbrace{z_m - \frac{n+\Delta n}{1+\frac{1}{e}}\left(\frac{z_m}{n+\Delta n}\right)^{1+\frac{1}{e}}}_{u_w^*} \tag{4}$$

Using the marginal wife buncher solution for the linear setting, $n+\Delta n = z_m + \Delta z_m$, we rearrange it to:

$$\frac{1}{1+\frac{\Delta z_w}{z_m}} - \frac{1}{1+\frac{1}{e}}\left[\frac{1}{1+\frac{\Delta z_w}{z_m}}\right]^{1+\frac{1}{e}} - \frac{1}{1+e}\left[1-\tau\right]^{1+e} = 0 \tag{5}$$

Where equation (5) characterizes the relationship between percentage relative earnings response $\frac{\Delta z_w}{z_m}$, the implicit male breadwinner tax τ and the earnings elasticity e. As the percentage relative earnings response is estimated through the excess mass bunching estimator $b = \frac{B}{h_0}$ at the notch point, this equation will be our empirical link between e and τ in the following sections.

4 Empirical analysis

4.1 Empirical model

To estimate the excess mass b discussed in the previous chapter, we need to construct a counterfactual distribution of the wife's relative earnings in the absence of the male breadwinner norm. I start with the following standard regression to estimate it through polynomial fitting:

$$c_j = \sum_{i=1}^{q} \beta_i \cdot \left(\theta^j\right)^i + \sum_{i=\theta_-}^{\theta_+} \delta_i \cdot 1\left[\theta^j = i\right] + \nu_i \tag{6}$$

Where c_j is the number of wife's income shares at the j bin, q is the polynomial order and θ^j the wife's relative income at bin j. The 'small bunching window' $[\theta_+, \theta_-]$ is the excluded range, chosen to include all of the bunching couples, and ν_i is the random error component. Therefore, the predicted count at the excluded range $\hat{c}_j$ will be given by:

$$\hat{c}_j = \sum_{i=1}^{q} \hat{\beta}_i \cdot (\theta)^i \tag{7}$$

Which we estimate iteratively by ordinary least squares following Chetty et al. (2011). The number of bunching couples will be then given by the sum of differences:

$$\hat{B} = \sum_{i=\theta_-}^{\theta_+} (c_j - \hat{c}_j) \tag{8}$$

Where finally we reach the normalized excess mass estimate by scaling through the counterfactual frequency in the excluded range:

$$\hat{b} = \frac{\hat{B}}{\sum_{i=\theta_-}^{\theta_+} \hat{c}_j} \tag{9}$$

This methodology has a few considerations that should be mentioned. The estimation above the notch may be quite sensitive to assumptions such as the order of the polynomial and the width of the excluded range. For this reason, our key results in the following sections will not rely on being able to estimate the counterfactual distribution and 'missing mass' after the notch. Our estimation of the tax rate associated with male breadwinner norm will only make use of the bunching estimates from this procedure.

The $\hat{b}$ should also be understood as a representation of intensive margin responses to the notched incentives. However by introducing a discrete jump in social cost over the threshold, the norm may also create extensive margin responses where women choose to leave the labor force and their share of earnings drop to zero. While this would change the counterfactual distribution, as discussed in Kleven and Waseem (2013), the estimation procedure intends to provide a "partial counterfactual" without intensive responses, but not extensive responses.

4.2 Data

We use the nationally representative Census survey in Brazil, the *Censo Demográfico*, assembled by the *Instituto Brasileiro de Geografia e Estatística* (IBGE) for the years 1991, 2000 and 2010. The Census survey collects data by interviewing all households across the country, and its main questionnaire consists of basic demographic information and population count. Concurrently with the main questionnaire, IBGE conducts a more detailed interview with a large random sample of households, in which they survey more detailed socio-demographic information from all residents, including total earnings, familial information and relationship status. This latter survey will be the basis for our empirical analysis. For our initial sample, we will use the universe of dual-earning couples where both spouses are within 25 and 65 years of age, cohabiting together, and the male is described as head of the household.

4.3 Wife's share of couple earnings revisited

Before reporting our bunching results, we'll utilize the 2010 census data to address some recent valuable discussion in the literature over the use of the wife's share of couple earnings for identification of the male breadwinner norm, and build over these points to increase robustness on our subsequent estimations.

Zinovyeva and Tverdostup (2021) suggests that the discontinuity at the 0.5 point of wife's relative

household earnings might not be because of gender norms, but instead may be driven by how couples split earnings when they have the same job or are self-employed. They support this view primarily by showing that the distribution discontinuity seems to disappear for a finnish administrative dataset, using a McCrary test, when excluding co-working and self-employed couples. Inversely, when looking exclusively into those couples, a sizeable discontinuity around the 0.5 thresold appears - many of them have equal or very similar earnings.

Zinovyeva and Tverdostup explain this difference by arguing that self-employed spouses want to split earnings equally for the sake of simplicity and/or tax reasons, and that couples who work for the same firm may be exposed to less informational asymmetry or job role differences, and therefore will have more similar pay relative to what they would get if working for two different employers.

I initially test if these patterns hold for the brazilian case by plotting those same groups using census data. Since the brazilian census doesn't have firm-level information, I instead categorize co-working status using 4-digit activity- and occupation-level codes for spouses within 25 and 65 years of age cohabiting together. Illustrative results for the 2010 survey are shown in Figure 3, with bin widths of 0.025. Although the frequency of couples at the midpoint bin is expectedly smaller for spouses in different firms or not self-employed, visually it doesn't seem that the discontinuity can be explained away by this hypothesis, as will be shown.

Figure 4: Relative earnings of women, by co-working status – Censo Demográfico 2010

a) Different firms, not self-employed **b)** Same firm or both self-employed

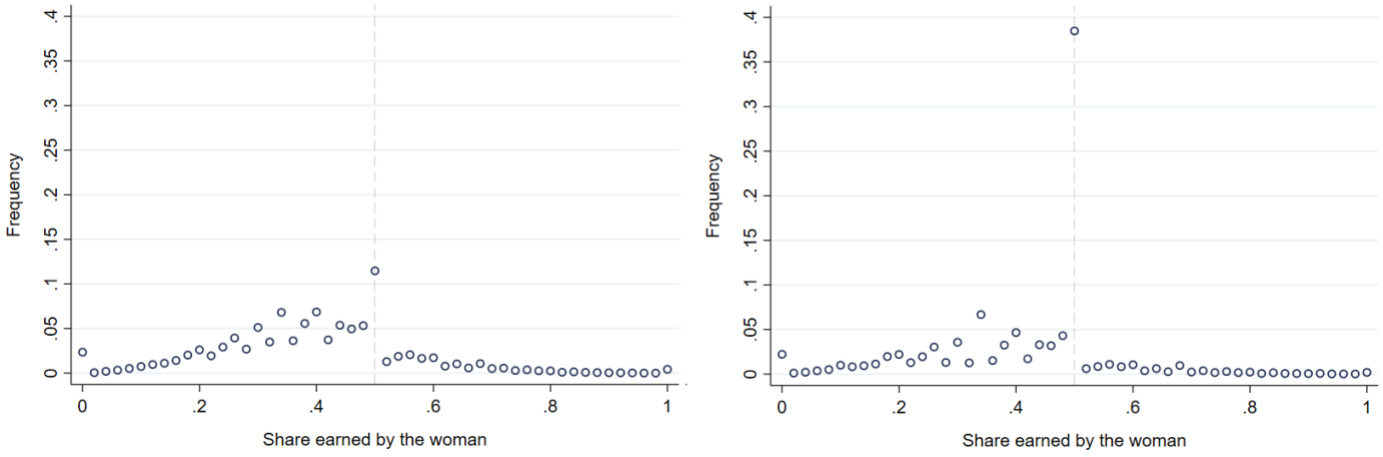

Similarly to the co-working hypothesis, a rigidity is introduced by minimum wage laws. As a fixed, lower-bound rate for the entire labor force, the minimum wage creates a higher frequency at the 0.5 relative earnings bin than what would otherwise be. Equal earning minimum wage couples may not exemplify a behavioral rectification of norm-breaking earning disparities, but instead a rigidity where working spouses who would make less than the given rate accumulate.

Another different source of identification issue is how open-ended survey questions about total

earnings may suffer response biases related to rounding effects (Schwabish, 2007; Roth and Slotwinski, 2020). Rounding occurs when survey respondents report total earnings equal to a close, arbitrary round number, erasing underlying earning gaps and affecting estimates sensitive to a specific threshold. As the male breadwinner norm is a relative measure, round number fixed effects cannot be used in the standard estimation, which creates identification problems. In Figure 4 we explore this issue by omitting earning multiples of 500 and 100 from the data. This exercise actually increases the fraction of very similar or equal earning couples at each new restriction, which we attribute to higher relative importance of the alternative issues discussed previously on the remaining samples created.

Figure 5: Relative earnings of women, by round number exclusion – Censo Demográfico 2010

(a) Total sample

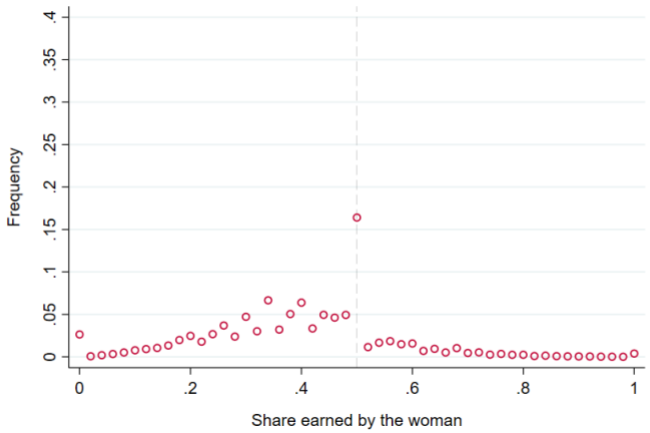

(b) 500 multiples removed

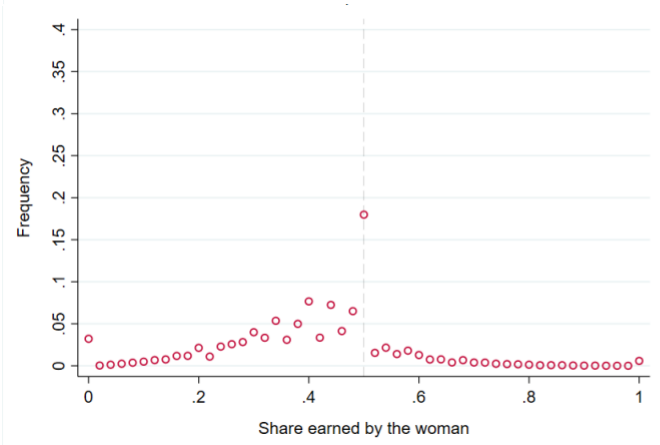

(c) 100 multiples removed

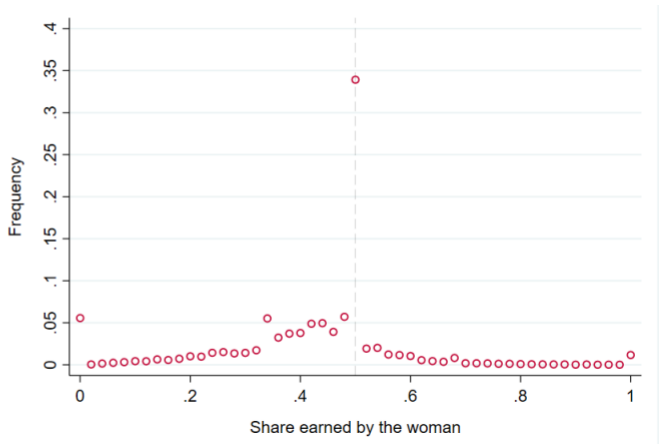

In order to assess the array of these mentioned alternative sources of discontinuity at the 0.5 relative earnings bin, we proceed by creating restricted samples that select for each of the issues, interact them, and then test for the discontinuity. First, we divide the sample by co-working and self-employed status; for rounding effects, we use both of the previous samples and add no de-rounding, de-rounding over

earning integers multiples of 500 and multiples 100; and, for minimum wage couples, we condition each preceding sample omitting and not omitting them. Then, we employ a McCrary test on each interaction.

McCrary (2008) developed an estimator for a density function discontinuity at a selected cutoff, where the test is implemented as a Wald test for the null hypothesis that the discontinuity is zero. The estimation consists of two local non-linear functions used to fit the density function on both sides of the threshold. Given a histogram of the relative earnings, the discontinuity estimator is given by:

$$\hat{\theta} = ln\hat{f}^{+} - ln\hat{f}^{-} \tag{10}$$

Where $\hat{f}^{+}$ is the fitted density estimate of the data points just after the threshold, and $\hat{f}^{-}$ just before. As in Bertrand et al. (2015), we chose the relative earnings threshold to be 0.501, and the results for each interaction can be found in Table 1. Log difference estimates suggest that the discontinuity persists even when considering all circumstances of co-working status, minimum wage coupling and 2-digit round number biases.

Table 1: McCrary density test for the 0.501 threshold

		Round numbers omitted		
		None	500 multiples	100 multiples
Non co-working, not self-employed	All sample	-1,579	-1,801	-2,105
		(0.006)	(0.008)	(0.010)
	MW omitted	-1,386	-1,294	-1.462
		(0.007)	(0.008)	(0.019)
Co-working or both self-employed	All sample	-3,454	-3,672	-3,828
		(0.015)	(0.022)	(0.030)
	MW omitted	-3,317	-3,415	-3,456
		(0.015)	(0.022)	(0.034)

This seemingly persistent discontinuity around the point where women start making more than their spouses, combined with unambiguous evidence of a male breadwinner norm in society at large, sets the groundwork for an investigation of possible behavioral earnings responses within couples to avoid threats to the breadwinner status of husbands. In subsequent sections, I'll make use of the restricted sample specifications of non-coworking, non self-employed, and non-minimum wage earning couples - with the appropriate earning multiples omitted for each census years - to undertake the non-parametric bunching estimations.

4.4 Assortative matching

An important consideration is made by Binder and Lam (2020), who approach the discontinuity at the halfway point of the wife's relative earnings distribution through the lens of standard marital matching models. They argue that a male breadwinner norm cannot be inferred from the data alone, as skewed distributions of spousal earnings can be generated through marriage markets that result in positive assortative matching and make no assumptions about underlying preferences.

Their argument is based on Becker (1974) and runs as follows: if utility is transferable between spouses and individuals are characterized by multiple attributes which include earnings, then the distribution of earnings gaps in a marriage market where the marriage output is a function of these attributes will clearly depend on how they're are correlated with earnings in the population. The beckerian matching equilibrium can consist of sorting on earnings, without any specific social norm-based preference, if the collection of attributes matching is positively correlated with earnings in the population, as some degree of sorting on earnings must occur after market clearing. And, given a significant gender gap in earnings, this model could predict a result in which a higher number of wives have similar or equal earnings as their husbands and fewer have higher earnings, regardless if the joint distributions could sustain a larger fraction of out-earning wives.

In general, what these models predict is that the rank-order, rather than the level of incomes, will tend to coincide between wives and husbands under minimal dependency assumptions in marriage formation, and as such the relative earnings distribution can easily be skewed unless the earnings distributions for both spouses are similar on level.

What implicitly follows this setting is that, if it is available to us the attributes underlying the matching process, we should arrive at a distribution of the wife's relative earnings with the same key patterns as the empirical one by randomly rearranging couples over these attributes. If such patterns cannot emerge, it's reasonable to assume that the matching process isn't the reason for the midway point discontinuity or the small fraction of wives out-earning their husbands.

We attempt to simulate these random matches using Brazilian data. First, we divided our main specification sample of husbands into attribute groups combining the following characteristics: age (three age groups: 25-39, 40–49 and 50–65 years old); educational level (four schooling levels: without schooling or incomplete primary education, complete primary education or incomplete high school, complete high school or incomplete undergraduate and complete undergraduate or higher); region of residence (five regions: North, Northeast, Southeast, South and Midwest); and religion (five religious groups: catholic, traditional evangelical, pentecostal evangelical, neopentecostal evangelical and irreligious or others). Then, we randomly re-matched husbands to wives within each collection of attribute groups, and checked the resulting kernel density distribution estimations. Figure 4 shows our results for this re-matching exercise through an increasing number of attribute groups.

In Figure 6.a, the rematch occurred only over age groups; 6.b over age · education; 6.c over age · education · region; and, finally, 6.d over all the preceding groups and religion. Although these are not

exhaustive attribute groups in which couples may choose to match, our simulation exercises show that a discontinuity point around the 0.5 relative earnings threshold doesn't seem to emerge, and a McCrary test over the simulated 4.d sample confirms so. These results give credence to the idea that the observed empirical discontinuity occurs after couples are together, which allows us to investigate further.

Figure 6: Kernel density of couple re-match over attribute groups – Censo Demográfico 2010

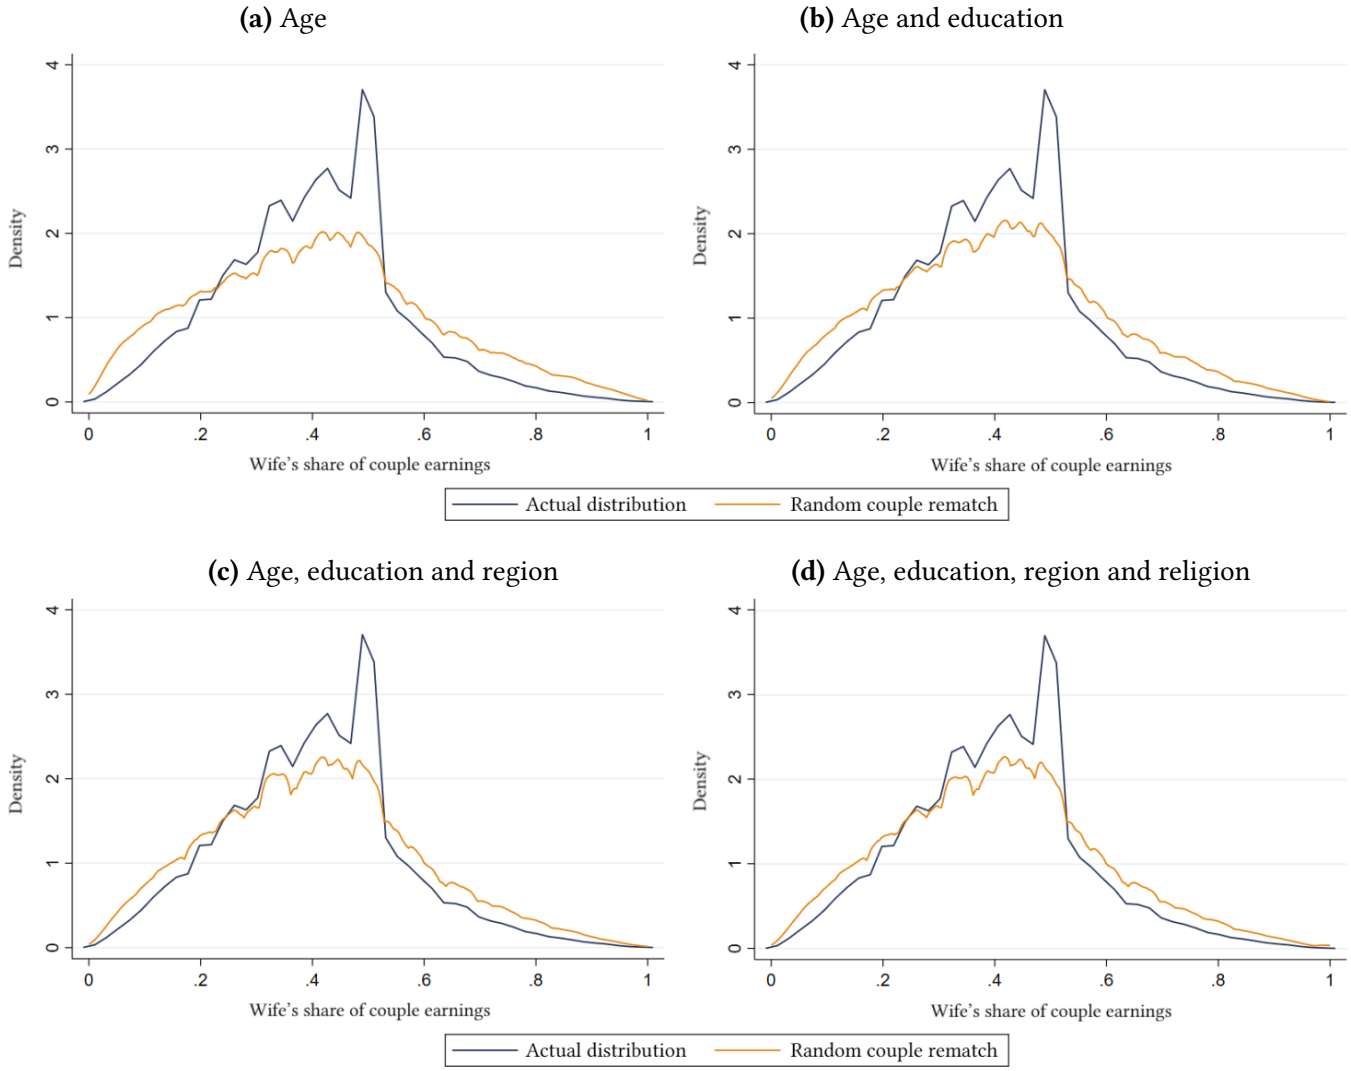

4.5 Bunching over time

This section presents our main bunching results. Figure 7 considers the total bunching behavior in our de-rounded sample of non-coworking, non self-employed and non-minimum wage couples in Brazil, where each panel represents a given census year and shows the empirical distribution of θ (blue line) as a histogram in 0.025 bins, along with the estimated 9th degree counterfactual distribution (red line).

Figure 7: Bunching results over time – Censo Demográfico

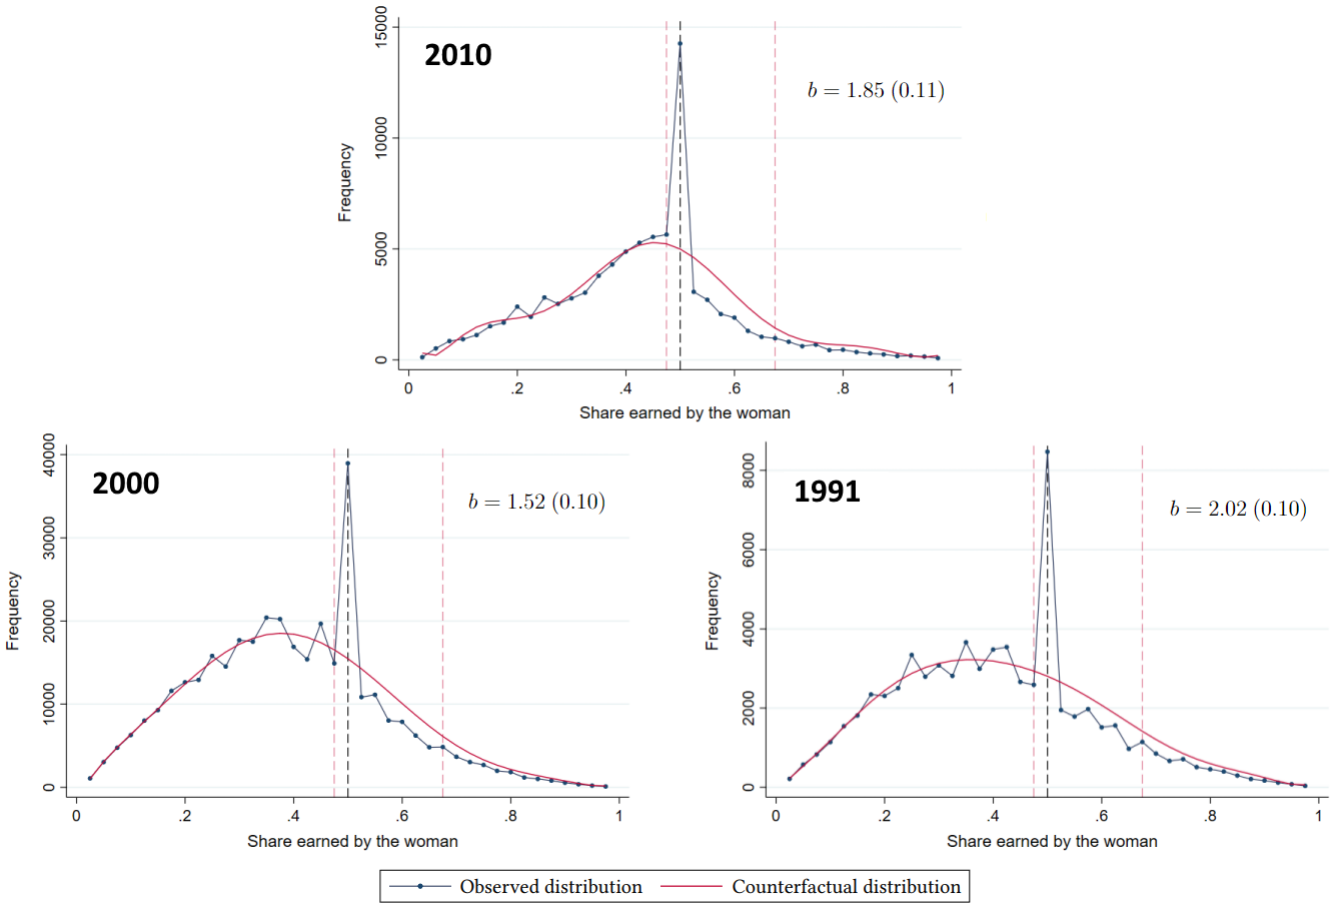

All standard errors are obtained by bootstrapping the procedure 200 times, where in each replication I draw from the empirical distribution of residuals with replacement, and generate a new vector of frequencies to obtain a new estimate of b. The standard error is then computed as the standard deviation of the resulting distribution of bs.

Some points about earnings de-rounding are worth noting. In the 2010 census survey, we chose to de-round over 100 multiples, however rounding effects associated with these multiples might not be the same for prior census years. This is especially relevant in 1991, where Brazil suffered from hyperinflation and the minimum wage was Cr$17.000,00. To counter these issues, we chose to de-round the 1991 and 2000 samples using multiples which appeared with similar frequency over the empirical bins across censuses. For instance, the notch bin was 73% composed of 100 multiples in 2010, and the most comparable round number multiple to this in the 2000 census was 50 (56%), and for 1991 was 2000 (68%), which we picked.

We proceed by using our normalized excess mass estimates from our sample to acquire a tax-like representation of the male breadwinner norm. In table 2, we assume a plethora of compensated labor

supply elasticities, particularly those in the range established by Mattos and Terra (2016), and, through equation (5), calculate the associated gender norm tax rate on women's earnings.

Our highest implicit gender norm tax rate occurs on the 1991 census sample, where it ranged from 20% of women's income to up to 43% given our selected range of elasticities. Those rates fell in the following censuses, with the lowest point on the 2000s. However, our bootstrapped standard deviations are relatively high, indicating high spread in the mean for each year.

Table 2: Implicit gender norm tax rate estimates over time

		1991		2000		2010	
q	e	b	τ	b	τ	b	τ
9	0.40	2.02 (0.10)	0.43 [0.23; 0.62]	1.52 (0.10)	0.35 [0.15; 0.53]	1.85 (0.11)	0.40 [0,18; 0,61]
	0.57		0.36 [0.16; 0.55]		0.29 [0.10; 0.48]		0.34 [0.12; 0.55]
	1.36		0.20 [0.09; 0.40]		0.15 [0; 0.34]		0.19 [0; 0.40]

4.6 Group heterogeneity

We'll now attempt to analyse how sub-groups within our main sample might behave heterogeneously with respect to the breadwinner norm. Two main categories will be analysed: religiosity and age. Figure 8 shows bunching behavior within the non-religious (or irreligious) and the religious sub-samples (the latter given by catholics, traditional evangelicals, pentecostal evangelicals and neopentecostal evangelicals).

Figure 8: Bunching results by religious affiliation – Censo Demográfico 2010

(a) Non-religious **(b)** Religious

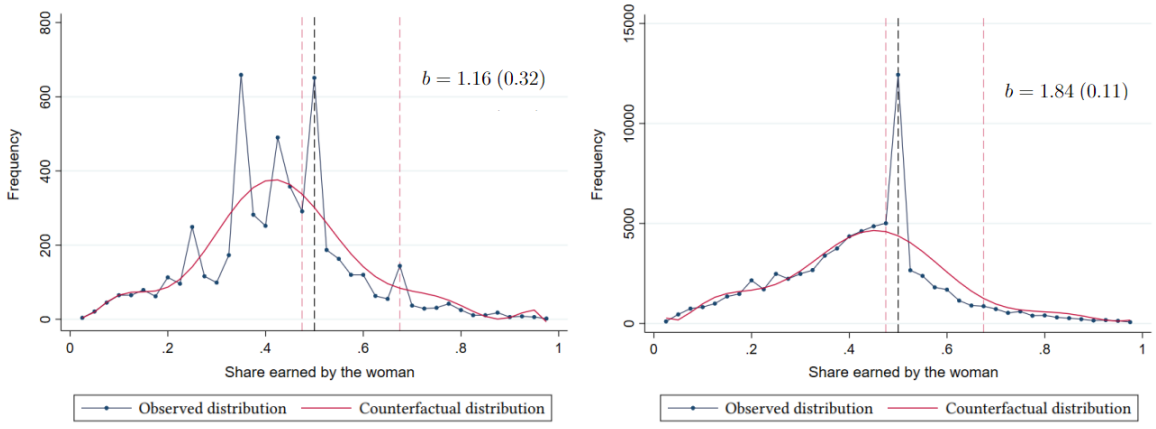

Non-religious couples bunching is shown to be smaller and have high bootstrapped standard deviations, whereas religious couples, which incorporates most of our sample, have more pronounced bunching and less spread. The same analysis over the age dimension is shown in Figure 9. Figure 9a includes couples where both are between 50 and 65 years of age, and Figure 9.b for couples between 25 to 49 years of age.

Figure 9: Bunching results by age – Censo Demográfico 2010

(a) 50-65 **(b)** 25-49

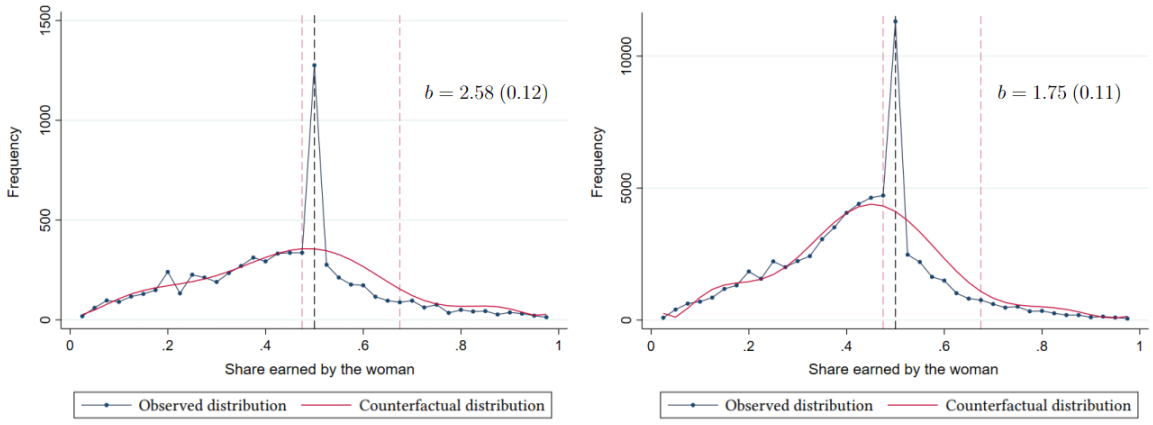

In Table 3 we have a summary of our group heterogeneity implicit tax estimates, as well as the mean net difference between the groups, which can be implied as a gender norm tax increase when considering religion adherence and/or old gender norm customs.

Table 3: Implicit gender norm tax rate by heterogeneous groups

Religious affiliation			Religious		Non-religious		Mean net
	q	e	b	τ_r	b	τ_{nr}	$\tau_r - \tau_{nr}$
	9	0.40	1.84 (0.11)	0.40 [0.19; 0.61]	1.16 (0.32)	0.28 [0; 0.9]	0.12
		0.57		0.34 [0.13; 0.55]		0.23 [0; 0.85]	0.11
		1.36		0.18 [0; 0.39]		0.12 [0; 0.74]	0.06
Age			50-65		25-49		Mean net
	q	e	b	τ_o	b	τ_n	$\tau_o - \tau_n$
	9	0.40	2.58 (0.12)	0.49 [0.26; 0.72]	1.75 (0.11)	0.39 [0.18; 0.6]	0.1
		0.57		0.42 [0.19; 0.65]		0.32 [0.11; 0.53]	0.1
		1.36		0.24 [0.03; 0.45]		0.21 [0; 0.42]	0.03

5 Conclusion

Gender norms' influence on real economic variables are tangible and can be, under certain contexts, quantitatively analysed. We argue that a male breadwinner norm can act as a tax on female earnings by explicitly incorporating the gender norm in modeling women's decision-making, and so providing a conceptual framework to better understand one of many factors that may limit women's earnings in their lifetimes.

We find that bunching behavior around the 0.5 threshold of women's share of couple earnings cannot be explained away by hypotheses related to co-working or self-employment behavior, round number effects, minimum wage bias or assortative matching. By using a range of labor supply elasticities, we employ our bunching estimates to ascribe implicit tax rates related to the male breadwinner norm. In doing so, we find ranges that can go considerably high, and adversely impact couples of more older and religious cohorts.

Our identification strategy can of course be criticized. Administrative, panel-level data could be used in the future for a more granular approach where alternative sources of discontinuity are causally considered over different time frames, methodologies and groups. It would also be important to gain a better understanding of the channels in which those earnings adjustments occur, considering the type

of social constraints applied to gender norm deviations and looking through the lens of both partners, as this will shed more light on how a seemingly welfare-reducing norm continues to exist.

Referências

AKERLOF, George; KRANTON, Rachel. Economics and Identity. **The Quarterly Journal of Economics**, 115(3), p. 715–753, 2000.

BECKER, Gary. A theory of social interactions. **Journal of Political Economy**, p. 1063–1093, 1974.

BERTRAND, Marianne; KAMENICA, Emir; PAN, Jessica. Gender Identity and Relative Income within Households. **The Quarterly Journal of Economics**, v. 130, p. 571–614, 2015.

BIANCHI, Suzanne et al. Is anyone doing the housework? Trends in the gender division of household labor. **Social forces**, v. 79.1, p. 191–228, 2000.

BINDER, Ariel; LAM, David. Is there a male breadwinner norm? The hazards of inferring preferences from marriage market outcomes. **National Bureau of Economic Research**, w24907, 2018.

BITTMAN, M et al. When does gender trump money? Bargaining and time in household work. **American Journal of Sociology**, v. 109.1, p. 186–214, 2003.

CASELLI, Francesca; WINGENDER, Philippe. Bunching at 3 percent: The Maastricht fiscal criterion and government deficits. **International Monetary Fund**, 2018.

CAVAPOZZI, Danilo; FRANCESCONI, Marco; NICOLETTI, Cheti. The Impact of Gender Role Norms on Mothers' Labor Supply. **Journal of Economic Behavior Organization**, v. 186, p. 113–134, 2021.

CHETTY, Raj; AL., et. Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records. **The Quarterly Journal of Economics**, p. 749–804, 2011.

CODAZZI, Karen; PERO, Valéria; SANT'ANNA, André. Social norms and female labor participation in Brazil. **Review of Development Economics**, v. 22.4, p. 1513–1535, 2018.

COLLINS, Matthew; CURTIS, John. Bunching of residential building energy performance certificates at threshold values. **Applied Energy**, v. 211, p. 662–676, 2018.

CONNELL, Raewyn. Gender. **Cambridge, UK: Polity Press**, 2010.

COOKE, Lynn. "Doing" gender in context: Household bargaining and risk of divorce in Germany and the United States. **American Journal of Sociology**, 112(2), p. 442–472, 2006.

DIAMOND, Rebecca; PERSSON, Petra. The long-term consequences of teacher discretion in grading of high-stakes. **National Bureau of Economic Research**, w22207, 2016.

DONGCHENG, Han; FANBO, Kong; ZIXUN, Wang. Gender identity and relative income within households: Evidence from China. **arXiv preprint**, 2021.

FOGLI, A.; VELDKAMP, L. Nature or nurture? Learning and the geography of female labor force participation. **Econometrica**, 79(4), p. 1103–1138, 2011.

GOURIO, F; ROYS, N. Size-dependent regulations, firm size distribution, and reallocation. **Quantitative Economics**, v. 5, p. 377–416, 2014.

GREENSTEIN, Theodore. Economic dependence, gender, and the division of labor in the home: A replication and extension. **Journal of Marriage and Family**, v. 62.2, p. 322–335, 2000.

GRUBB, M; OSBORNE, M. Cellular service demand: biased beliefs, learning, and bill shock. **American Economic Review**, v. 105, p. 234–71, 2015.

_____. Labor supply and taxes: A survey. **Journal of Economic Literature**, p. 961–1075, 2011.

KILLEWALD, Alexandra. Money, Work, and Marital Stability: Assessing Change in the Gendered Determinants of Divorce. **American Sociological Review**, 81(4), p. 696–719, 2016.

KLEVEN, Henrik; WASEEM, Mazhar. Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan. **The Quarterly Journal of Economics**, v. 128.2, p. 669–723, 2013.

KLEVEN, Henrik Jacobsen. Bunching. **Annual Review of Economics**, v. 8, p. 435–64, 2016.

MATTOS, Enlinson; TERRA, Rafael. Nature of transfers, income tax function and empirical estimation of elasticity of taxable income for Brazil. **Applied Economics**, 48(53), p. 1–20, 2016.

MCCRARY, Justin. Manipulation of the running variable in the regression discontinuity design: A density test. **Journal of Econometrics**, v. 142.2, p. 698–714, 2008.

RISMAN, Barbara. Gender Vertigo: American Families in Transition. **Yale University Press**, 1999.

SAEZ, Emmanuel. Do Taxpayers Bunch at Kink Points? **American Economic Journal: Economic Policy**, p. 180–212, 2010.

SLOTWINSK, Michaela; ROTH, Anja. Gender norms and income misreporting within households. **CESifo Working Paper**, p. 2364–1428, 2021.

SPRENGHOLZ, Maximilian; WIEBER, Anna; HOLST, Elke. Gender identity and wives' labor market outcomes in West and East Germany between 1983 and 2016. **Socio-Economic Review**, v. 20.1, p. 257–279, 2020.

SURVEY, World Values. WVS: Round Seven – Country-Pooled Datafile. **WVSA Secretariat**, 2020.

ZINOVYEVA, Natalia; TVERDOSTUP, Maryna. Gender identity, coworking spouses, and relative income within households. **American Economic Journal: Applied Economics**, v. 13.4, p. 258–84, 2021.