

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/141183>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Gender Identity, Co-Working Spouses and Relative Income Within Households

Natalia Zinovyeva and Maryna Tverdostup*

September 24, 2020

Forthcoming in American Economic Journal: Applied Economics

Abstract

Bertrand, Kamenica and Pan (2015) document that in the United States there is a discontinuity to the right of 0.5 in the distribution of households according to the female share of total earnings which they attribute to the existence of a gender identity norm. We provide an alternative explanation for this discontinuity. Using linked employer-employee data from Finland, we show that the discontinuity emerges as a result of equalization and convergence of earnings in co-working couples, and it is associated with an increase in the relative earnings of women, rather than a decrease as predicted by the norm.

Keywords: co-working spouses, gender identity norms, spouses' relative earnings

JEL Classification: D10, J16, J21

According to the World Values Survey (Inglehart et al., 2014), 36% of Americans agree with the statement “[I]f a woman earns more money than her husband, it’s almost certain to cause problems.” In the European Union, this view is shared by 39% of the population. Bertrand, Kamenica and Pan (2015) suggest that a social norm assigning a bread-winner role to the husband may affect the formation and separation of couples and it can induce high-earning married women to reduce their labor supply. Using

*Zinovyeva: University of Warwick, Coventry, CV4 7AL, United Kingdom, and Aalto University, PO Box 21210 FI-00076 AALTO (e-mail: natalia.zinovyeva@warwick.ac.uk); and Tverdostup: University of Innsbruck, Universitätsstrasse 15, Innsbruck, Austria, and University of Tartu, Estonia (e-mail: maryna.tverdostup@uibk.ac.at). We would like to thank Daron Acemoglu, Joshua Angrist, David Autor, Manuel Bagues, Clare Balboni, Pamela Campa, Daniel Deming, Amy Finkelstein, Emily Nix, Eric Plug, Debraj Ray, Matti Sarvimäki, Sarah Smith, Marko Terviö, Mikhail Zinovyev, Alexandre Mas (editor) and three anonymous referees, as well as participants of presentations at Clark University, USC Marshall, MIT, University of Essex, University of Bristol, University of Warwick, University of Innsbruck, SITE, VATT, Tinbergen Institute, IZA World Labor conference, and Nordic Summer Institute in Uppsala for their useful comments and suggestions. We acknowledge financial support from the Yrjö Jahnsson Foundation, HSE Foundation, Dora Plus 1.2 Grant, and Archimedes Foundation. Some results in this paper were obtained from a validation analysis conducted by Census Bureau staff using the SIPP Completed Gold Standard Files (U.S. Census Bureau, 2015a) and the programs written by the authors. We are particularly grateful to Jordan C. Stanley for his assistance with this process. The programs were originally run on the SIPP Synthetic Beta (U.S. Census Bureau, 2015b) available on the Synthetic Data Server at Cornell University funded by NSF Grant SES-1042181. The validation analysis does not imply endorsement by the Census Bureau of any methods, results, opinions, or views presented in this paper.

data from the United States, they argue that this gender identity norm generates a discontinuity to the right of 0.5 in the distribution of married couples by the share of income earned by the wife. The discontinuity can be observed already in recently married couples and it grows with marriage tenure. Bertrand, Kamenica and Pan (2015) point out that it cannot be explained by classical marriage market theories. Models that consider marriage as a partnership for the purpose of joint production and joint consumption do not predict anything unusual around the point where spouses have similar earnings. Similarly, models that consider marriage as a source of gains from specialization do not attribute any particular significance to the 0.5 point.

The existence of a discontinuity to the right of 0.5 in the relative earnings distribution has been widely cited both in the media and in academia as evidence for the relevance of the gender identity norm. Some authors have also pointed out that a substantial part of the discontinuity is due to the existence of a point mass of couples exactly at 0.5 (Hederos and Stenberg, 2015, Binder and Lam, 2018).¹ As shown in panel (a) of Figure 1, the discontinuity to the right of 0.5 estimated by Bertrand, Kamenica and Pan (2015) becomes smaller if spouses with equal earnings are excluded, with the McCrary (2008) estimate dropping from 12.3% to 7.4%.²

In this paper, we provide evidence contradicting the social norm interpretation of the discontinuity (and the point mass) at 0.5 and we propose an alternative explanation. We use linked employer-employee data from Finland that has detailed information on the individual employment and earnings history of the entire population of Finnish individuals for the period between 1988 and 2014. Although women in Finland have achieved a relatively high degree of equality in many dimensions, survey information suggests that the gender norm regarding relative earnings in the households is as relevant as in the United States. According to the World Values Survey 1995-1998 (Inglehart et al., 2014), 33.9% of Finns agreed with the idea that a woman should earn less than her husband to avoid family problems.

As in Bertrand, Kamenica and Pan (2015), in Finland, we observe a drop at 0.5 in the distribution of households according to the female share of total earnings (see panel (b) of Figure 1). The estimate of the discontinuity is of a similar magnitude as in the United States and, similarly, about a third of the drop is due to the excess mass of spouses with identical earnings. The discontinuity is also present in the first year of marriage. However, we document several additional facts that, jointly, suggest that the gender identity norm is not the cause for the observed shape of the distribution around 0.5.

First, we examine the distribution of relative earnings at the beginning of cohabitation, which provides a better proxy of the time of union formation than marriage. We find no significant discontinuity at this stage of the relationship, suggesting that the gender norm does not affect the formation of couples in a discontinuous way.

Second, the norm does not seem to play a role for separations either. Separation

¹Hederos and Stenberg (2015) show that in Sweden, the discontinuity can be partly attributed to the existence of an excess mass of couples with identical earnings, particularly, in couples where at least one of the spouses is self-employed. Binder and Lam (2018) point out that also in the United States, the discontinuity is primarily driven by an excess mass of couples with identical earnings.

²Following Bertrand, Kamenica and Pan (2015), we use administrative information on earnings in married couples from the U.S. Census Bureau's Survey of Income and Program Participation linked to information from the Social Security Administration and the Internal Revenue Service (U.S. Census Bureau, 2015a). We use SIPP Completed Gold Standard Files, in which some observations are substituted with imputed values. We drop these observations.




Figure 1: Relative earnings of women

Notes: The data on panel (a) are on married couples observed in SIPP/SSA/IRS Completed Gold Standard Files, 1990-2004. On panel (b), the data is on cohabiting and married couples from FLEED, 1988-2014. In both cases, the sample is restricted to couples with both partners aged between 18 and 65 years and receiving positive earned income. Each dot indicates a fraction of couples in a 2% relative income bin; bins are right-closed. The dashed line is the lowess smoother applied to the distribution allowing for a break at 0.5. The dark-colored crosses and dashed line show the fraction of couples in each bin and the lowess smoother calculated after excluding households with identical earnings of both spouses.

rates do not exhibit any discontinuity around the 0.5 threshold of relative earnings. Instead, the relationship between the probability of separation and the relative earnings distribution exhibits a U-shape, with higher separation rates among couples with large earnings differentials either in favor of the husband or in favor of the wife.

Third, the discontinuity in the distribution only arises in couples where both spouses are self-employed (around 6% of all employed couples) or work together in the same firm (around 9%). Hereafter, we refer to these two groups as co-working couples. For the rest of the population, there is no evidence of any unusual phenomena in the vicinity of the 0.5 point. The pattern looks different for these two groups of co-working couples. In the case of self-employed couples, the discontinuity to the right of 0.5 is mainly due to a significant fraction of couples bunching exactly at 0.5, while among spouses working for the same employer, the distribution exhibits a cliff at 0.5 with only a small fraction of couples having identical earnings.

Fourth, the observed dynamics rules out a more specific formulation of the gender identity norm theory, according to which the norm is activated only when spouses are jointly self-employed or work in the same firm. Theoretically, this may occur if co-working makes the comparison between spouses more salient or if adjustments in accordance with the norm are feasible only in self-employed couples. We find that the discontinuity does not arise as a result of a reduction in the share of couples where women slightly outearn their husbands, as the gender identity norm would predict. Instead, when couples on both sides of the distribution become self-employed, they tend to equalize earnings leading to an excess mass at 0.5. Similarly, when couples start working together in the same firm, there is a compression of earnings toward 0.5. Since initially there are more couples where women earn less than men, this earnings compression creates a larger mass of couples just to the left of 0.5 than to the right of this point, which statistical tests identify as a discontinuity. Moreover, we also observe that co-working leads to an increase in female earnings above the earnings of similar women in non-co-working couples.

Overall, our results contradict the idea that the gender identity norm exhibits a discontinuity at the point of equal earnings. Some couples may prefer that the husband earns more than his wife, but small variations around the 0.5 point do not seem to make that much of a difference.

There are several factors that may generate earnings equalization in self-employed couples. Self-employed spouses may report identical earnings when they face strong fiscal incentives to do so.³ In Finland, such incentives may be generated by individual income tax filing under a progressive scale. In other countries where joint filing is possible, incentives for reporting equal earnings may be created by the existence of itemized deductions (e.g., deductions for mortgage interest on a loan or for medical expenses in the United States, the United Kingdom, and Germany). The legal environment may also incentivize equal income sharing. In some countries, including Finland, businesses registered as partnerships must split entrepreneurial profits equally between the partners in the absence of a specific written agreement. In fact, in our data we observe that the share of equal earners is the largest in partnerships, signaling the importance of legal defaults for income splitting. Many self-employed couples may also use rules of thumb for income splitting within the family. Consistently, we observe substantial

³There is evidence that self-employed couples may misreport their earnings to minimize the tax burden (Stephens and Ward-Batts, 2004, Schuetze, 2006, LaLumia, 2008, Kleven, Kreiner and Saez, 2009, Kabatek, van Soest and Stancanelli, 2014, Harju and Matikka, 2016).

bunching of relative earnings also at other salient shares, such as 0.4 and 0.3.

There are also several possible forces that may lead to earnings convergence in couples where spouses work in the same firm. Rent sharing and within-firm pay equity constraints may result in a lower within-firm variability in pay (Breza, Kaur and Shamdani, 2018, Dube, Giuliano and Leonard, 2019, Saez, Schoefer and Seim, 2019). To understand the role played by such earnings compression within firms, we randomly paired unrelated male and female co-workers and examined the evolution of the relative earnings distribution in these fictitious couples. We observe that, over the time the two randomly matched individuals work together in the same firm, their earnings converge and a discontinuity at 0.5 in the distribution of the female share arises. This result indicates that the earnings convergence and discontinuity observed among actual couples is, at least partly, induced by firm-level pay compression. In addition to firm-level dynamics, spouses co-working in the same firm may experience earnings convergence as a result of joint negotiations with the employer or, perhaps, direct favoritism. These mechanisms would explain why women in co-working couples earn more than observationally similar women in non-co-working couples.

Finally, we provide some suggestive evidence about the potential relevance of this mechanism for the U.S. labor market. As in Bertrand, Kamenica and Pan (2015), we use administrative information on earnings for a sample of U.S. households from 1990-2004. Unfortunately, the dataset does not offer information on whether spouses work in the same firm or are jointly self-employed. Instead, we use the available information on spouses' industry and occupation to identify a group of spouses with a larger probability of working together. As shown by Hyatt (2015), the likelihood to have a shared workplace tends to be particularly high among spouses working in the same industry and occupation. We find that, in this group of couples, the "missing" mass of households in which women slightly outearn their husbands is twice as large as in the overall population. Arguably, this pattern is more consistent with the hypothesis of earnings convergence in couples working together than with the explanation based on the existence of the gender identity norm.

The rest of the paper is organized as follows. Section 1 briefly illustrates how earnings convergence and equalization within couples can generate a discontinuity in the distribution of couples according to the wife's share of earned income. This section also formulates the main testable implications. In section 2, we describe the data used in the empirical part. In section 3, we document the existence of a discontinuity in the distribution of the female share in Finland, we explore whether the discontinuity is limited to co-working couples and analyze how the distribution changes around the time when spouses start working together. Finally, in section 4, we discuss the implications of our results.

1 Potential explanations of the discontinuity

Bertrand, Kamenica and Pan (2015) document that the distribution of households according to the relative earnings of women exhibits a discontinuity to the right of 0.5, which partly reflects the existence of a point mass at 0.5. Below, we first discuss how the gender identity norm may potentially generate this pattern. We then examine a number of alternative mechanisms that may generate empirically similar shape of the distribution. In particular, we study how the distribution of relative earnings within couples would be affected if (i) in some couples earnings were equalized, (ii) some

couples experienced convergence in earnings, and (iii) data on individual earnings were subject to rounding. We then outline several testable implications that allow us to distinguish among these theories empirically.

1.1 Gender identity norm

The gender identity norm in which women should not outearn their husbands may generate the distribution of relative earnings that is observed in the data in several ways. First, relatively fewer couples may form where women outearn men. Second, the likelihood of divorce may be higher in couples where women earn more than their husbands. Third, when couples are formed, high-earning women may adjust their labor supply so that their earnings do not surpass their husbands'. These adjustments may create a discontinuity at 0.5 in the distribution of households according to the relative earnings of women and, if individuals can perfectly manipulate their relative incomes, they may also create an excess mass of couples with spouses having identical earnings.

To illustrate a possible impact of the norm as well as the following arguments, we perform a simulation exercise. Female and male earnings are drawn from two gamma distributions where men are assumed to have higher average earnings. We construct couples by randomly matching men and women. The resulting distribution of the wife's share of earnings is shown in blue-colored dots in panel (a) of Figure 2.⁴

To simulate the impact of the gender identity norm, we consider a simple scenario where the norm affects a given proportion of couples where women initially outearn their husbands. To comply with the norm, these couples adjust the female share down until it reaches the point with equal earnings. Red dots on panel (a) show the relative earnings distribution after this adjustment takes place. A red line shows the estimate of the density function obtained using the McCrary (2008) procedure, allowing for discontinuity at 0.5. As in Bertrand, Kamenica and Pan (2015), we use right-closed bins. There is a visible drop in the distribution at 0.5. The McCrary test identifies a discontinuity also after if we exclude the 0.5 mass from the estimation, reflecting the decrease of the density estimate to the right of this point.


1.2 Equalization of earnings

Let us consider another mechanism that might affect the distribution of relative earnings. Some spouses may equalize earnings for other reasons than the gender norm, for instance, due to fiscal or legal incentives. Another source of earnings equalization may be the existence of salary schedules within firms. Many firms offer salary schedules to their workers that include a limited number of pay grades, increasing the probability that couples working in the same firm receive the same salary.

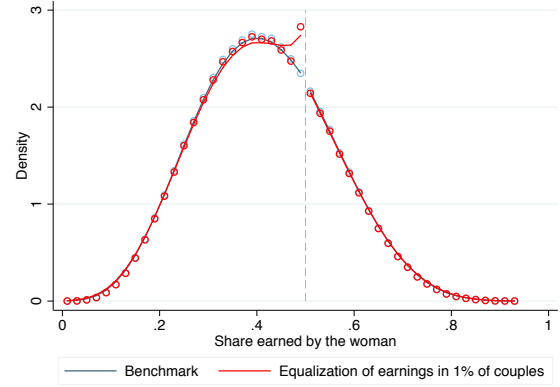
We illustrate this case in panel (b) of Figure 2. In this figure, a random proportion of couples on both sides of the distribution equalize earnings. The bunching at 0.5 emerges, creating an apparent discontinuity to the right of this point. However, note that no discontinuity would be observed if the mass at 0.5 was excluded from the distribution.

⁴In the figure, female and male earnings are distributed respectively as $\Gamma(5, 5000)$ and $\Gamma(7, 5000)$. Gamma-distribution and parameters of the distribution are selected to mimic the shape of the actual earnings distribution with no negative values, a large mass of individuals with moderate earnings, and a thin right tail. The average earnings of men and women in the simulated data are 35,000 and 25,000, respectively, resembling actual values of annual earnings in euros.


(a) 5% of couples to the right of 0.5 reduce female share exactly to 0.5




(b) 1% of couples equalize earnings



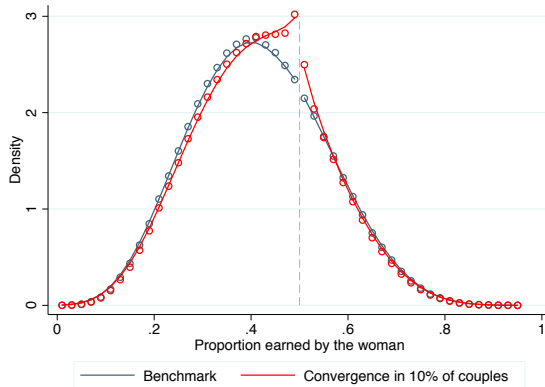
(c) Convergence by a 3-percentage point step toward 0.5



(d) Convergence by a random factor uniformly distributed between 0 and 1



(e) Convergence by a random factor in 10% of couples



(f) Rounded earnings

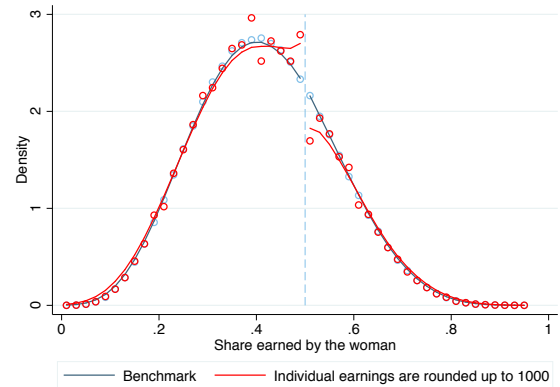


Figure 2: Relative earnings of women after various hypothetical adjustments, simulation

Notes: The figure uses simulated data to demonstrate how various forces discussed in section 1 can transform a smooth distribution of the relative earnings of women (shown in blue) into a distribution that exhibits a discontinuity at 0.5 (shown in red). To construct the data, we first assumed that female and male earnings are distributed respectively as $\Gamma(5, 5000)$ and $\Gamma(7, 5000)$. We then defined couples by randomly matching one million men and women.

The dots indicate a fraction of couples in a 2% relative income bin; bins are right-closed. The lines show the estimate of the density function obtained using the McCrary (2008) procedure with default estimation options, allowing for discontinuity just to the right of 0.5.

1.3 Earnings convergence within the couple

Some couples who work together may experience convergence in earnings as a result of rent sharing and fairness considerations at the firm level. As we explain below, this process may also generate an excess mass at 0.5 and a discontinuity to the right of this point.

We examine two types of convergence. First, we consider steplike earnings convergence. Couples on both sides of the distribution adjust their relative earnings by a certain fixed amount toward equal earnings. As shown in panel (c) of Figure 2, this adjustment compresses the distribution around 0.5 and creates a bunching at 0.5. Since initially there are more couples in which husbands outearn their wives than couples in which women earn relatively more, the process creates a discontinuity between the left and right limits of the density function, which persists when the mass at point 0.5 is excluded.

Second, we consider the case when the earnings of spouses converge by a random factor uniformly distributed between 0 and 1 (see panel (d) of Figure 2). In this case, earnings compression leads to an asymmetric peak distribution of relative earnings. Technically, the distribution is continuous with a kink at 0.5, however, the McCrary test produces a significant estimate of the drop at 0.5.⁵ In panel (e), we show the resulting distribution if only 10% of couples co-work and experience earnings convergence by a random factor. The distribution appears much smoother, but the discontinuity is still significant. Notably, the discontinuity shown in panels (d) and (e) is not related to the presence of an excess mass at 0.5. In fact, in the simulation presented here, no couples have identical earnings.

A common feature of these two processes is that changes in the rank order of individual earnings within the couple are not allowed. Earnings compression without restrictions on the rank order may not lead to a kink or a discontinuity.

1.4 Rounding of individual earnings

A related phenomenon is rounding. Very often the available data on individual earnings is subject to rounding, either by statistical offices or, in survey data, by respondents. To examine this issue, we round the individual earnings of individuals and compute the wife's share of earnings based on the rounded values. As we show in panel (f) of Figure 2, rounding generates several bunching points and discontinuities in the distribution of relative earnings, particularly to the right of 0.5.

This illustration shows that, even when the underlying distribution does not exhibit a discontinuity, rounding the data might generate one. Survey data might thus be inappropriate to study the existence of a discontinuity in the distribution of relative earnings. It might also be non-trivial to infer the existence of a gender identity norm in reporting based on the comparison of survey and administrative data.⁶

⁵The estimate is robust to the reduction of the default bandwidth to half, following the recommendations by McCrary (2008) for robust asymptotic inference.

⁶Murray-Close and Heggeness (2018) compare the earnings reported by participants in the Current Population Survey with their actual earnings, which they observe in administrative data. They find that in couples where women earn more than husbands, women are relatively more likely to underreport their earnings and men to overreport them, and they attribute this discrepancy to the impact of the gender identity norm on reporting.

1.5 Testable implications

While both the gender identity norm and earnings equalization or convergence in co-working couples may lead to an excess mass and a discontinuity to the right of 0.5 in the distribution of the female share of earnings, the two hypotheses have different implications in several other dimensions.

Start of the relationship The two hypotheses have different implications regarding the moment in the relationship when the discontinuity emerges. According to the gender identity norm hypothesis, the discontinuity may exist already at the start of the relationship, because of selective couple formation. Instead, the hypothesis of earnings convergence and equalization in co-working couples does not predict a discontinuity in the sample of newly formed couples, unless they were already working together.

Separation and divorce The gender identity norm hypothesis predicts that couples where women initially earn just a bit more than their husbands may have a higher likelihood to separate, either because they failed to anticipate the importance of the gender identity norm or because they only learned over time about the earnings of their partners. Instead, the forces that lead to earnings equalization or convergence are not expected to create a discontinuity or a kink in the separation rate at the point where spouses have identical earnings at the beginning of the relationship.

Affected couples The two hypotheses differ in their predictions regarding the type of couples that are expected to exhibit a discontinuity and an excess mass at 0.5. The gender identity norm should be more relevant among couples with more traditional values. Instead, the earnings convergence and equalization predict the discontinuity and a mass of equal earners only among couples that work together. In particular, spouses in self-employed couples should be more likely to earn the same, while spouses co-working in the same firm are expected to exhibit a compressed distribution of relative earnings around the point with equal earnings and a cliff at 0.5.

There is a refinement of the gender identity norm theory that might potentially rationalize the existence of a discontinuity uniquely in co-working couples. The norm itself may be “activated” only when both spouses work together because it makes the comparison of their positions and earnings more salient. It is also possible that only jointly self-employed couples may effectively implement the adjustments necessary to abide by the gender norm. Next, we consider several testable implications that allow disentangling this specific version of the gender identity norm theory from the earnings equalization and convergence hypothesis.

Dynamics The two hypotheses have different predictions regarding the changes in the relative earnings distribution around the start of co-working. The gender identity norm predicts that there will be a decrease in the share of couples in which the wife outearns her husband. Instead, the hypothesis of earnings equalization and convergence predicts that individual earnings in couples on both sides of the distribution will tend to equalize.

Moreover, for couples formed at the workplace, the two hypotheses have also a distinct prediction regarding the timing of the adjustments. The gender identity norm would predict that couples should reduce the female share soon after the start of

the relationship. In the case of the earnings convergence hypothesis, the prediction is ambiguous. If earnings convergence is due to firm-level forces that generate wage compression, the adjustment may already happen before the couple has been formed. Instead, if the adjustment reflects joint bargaining or information sharing within the couple, convergence may also arise after the start of the relationship.

Counterfactual earnings If co-working “activates” the gender identity norm, women who start co-working with their spouses are expected to earn, on average, less than comparable non-co-working women. The earnings convergence hypothesis does not provide a clear prediction. On the one hand, it is possible that co-working spouses are willing to accept a lower salary because they appreciate the joint location or other amenities associated with co-working. On the other hand, co-working spouses may have a better bargaining position with their employer or they may have access to better information allowing them to negotiate higher salaries.

Earnings distribution in placebo co-working couples The social norm assigning family-level gender roles provides no prediction regarding the relative earnings of unrelated employees working in the same firm. In particular, we should not expect to find a discontinuity in the distribution of the relative earnings of individuals in fictitious couples formed by randomly matching unrelated co-working women and men. Instead, the hypothesis of firm-level wage convergence does predict compression of the distribution of the female share and possibly also the emergence of a discontinuity.

2 Data

We use the Finnish Longitudinal Employer-Employee Data (FLEED), which contains registry information on the entire Finnish population from 1988 to 2014 (Statistics Finland, 2019). We restricted the sample to working-age individuals (18 to 65 years old). In this section, we describe the primary features of the dataset.

2.1 Couples

We follow the classification of Statistics Finland, which considers two individuals as a couple if they are cohabiting, married, or have a registered civil partnership.⁷ The sample includes around 2.6 million couples and each couple is observed on average for 11 years. About 1.5 million of these couples were formed after 1988, which is the first year available in our database.

The main variable of interest is annual labor earnings, which includes individual earned income and entrepreneurial income. Annual earnings are not top-coded and they are not rounded. As shown in Figure A1 in the Online Appendix, at the time of couple formation around 42% of a couple’s earnings are accrued by the woman. The relative earnings of women tend to rapidly decrease afterward, falling to 35% after six years into the relationship, and they slowly catch up later.

⁷Two individuals are considered cohabiting if they are of different sex, live permanently in the same dwelling, are at least 18 years old, their age difference is at most 15 years, they do not have a spouse and they are not siblings. 83% of couples are identified by Statistics Finland based on individuals’ cohabitation status.

2.2 Co-working spouses

Statistics Finland provides information on the identity of employers for about 90% of couples of wage-earners. Both spouses share the same employer in 9% of these couples and, within this group, two-thirds of couples work in the same establishment. Employer identifiers are not available whenever this information may allow for the identification of an individual, something that is more likely to happen in small family businesses. Therefore, the figure above may slightly underestimate the overall share of couples who work together. Furthermore, in 6% of couples, both spouses are self-employed. In general, we do not observe whether self-employed couples work together, but we observe that most of them work in the same 5-digit industry, suggesting that in most cases self-employed couples actually work together.

The workplace appears to be the meeting place for some spouses, however, for most couples, cohabitation precedes co-working. Around one-third of couples who ever co-work in the same firm were already working together when they started to cohabit. The median time to co-working in the same firm after cohabitation is 3 years and the median time to simultaneous self-employment is 6 years.

Women are more likely than men to change employer or main activity to become co-employed with their spouses. In couples where both spouses were already employed before co-working, the woman joins the firm of her husband in 47% of cases, the man joins the firm of his wife in 35% of cases, and in 18% of cases, both spouses change the employer. About 39% of women who start working with their partners were not employed the year earlier, while this is true only for 18% of men.

As shown in Figure A2 in the Online Appendix, there are no large differences in couples' initial relative earnings depending on whether they will co-work in the future or not. Women earn 44% of household earnings at the beginning of the relationship in couples who never co-work, compared to 43% in couples that will work the same firm and 41% if couples that become jointly self-employed.

3 Empirical Analysis

We start the analysis by documenting the existence of a sharp drop at the 0.5 point in the distribution of Finnish households according to the wife's share of earned income. We then use the testable implications outlined in section 1.5 to study whether this discontinuity is due to the gender norm or, alternatively, to earnings convergence and equalization among co-working couples. Finally, after showing that earnings equalization and convergence in co-working couples drive the discontinuity in Finland, we discuss the relevance of this mechanism for explaining the evidence from the United States.

3.1 Discontinuity in the relative income distribution

Panel (b) of Figure 1 shows the distribution of couples according to the wife's share of household labor earnings in a sample of couples where both spouses are employed and have positive earnings. On the y -axis, the figure reports the fraction of couples in a 2% relative income bins. As in Bertrand, Kamenica and Pan (2015), we use right-closed bins. About a third of this drop is due to the existence of an excess mass of couples where both spouses have the same earnings, who constitute 0.9% of all

employed couples. When these couples are excluded from the sample, the remaining drop at 0.5 is equal to 11.3% according to the McCrary test (Table 1, first row). The magnitude of the discontinuity is similar to the one observed in the United States.

As pointed out by Binder and Lam (2018), it is misleading to describe the excess mass at 0.5 as a “discontinuity” to the right (or to the left) of this point. Bunching at equal earnings creates a jump in the cumulative distribution of relative earnings and makes the theoretical density function at 0.5 equal to infinity. To avoid this problem, in what follows we always drop equal earners when we estimate the limits of the density function at 0.5.

3.2 Relative earnings in newly formed couples

Bertrand, Kamenica and Pan (2015) use the time of marriage as a proxy for the time of couple formation and show that a discontinuity exists among newlywed couples. In Finland, we also observe a significant drop in the estimated density function of relative earnings when we limit the sample to newly married couples. In this sample, 0.2% of spouses have identical earnings and a discontinuity is estimated to be 4.1%, significant at 1% (see Table A1 in the Online Appendix). However, this result is entirely driven by co-working couples, who constitute about 11% of employed newlywed couples. Among co-working couples, 2% have equal earnings and the discontinuity is 19.3%, while among non-co-working couples, there are no equal earners and the discontinuity is 2.9%, statistically insignificant.

In our data, we can also use the start of cohabitation as a proxy for the time of couple formation. In Finland, the median cohabitation tenure among newlyweds is 3.5 years and only 31% of initially cohabiting couples are observed to get married eventually. When we consider the earnings of spouses at the time when individuals start to cohabit, the distribution does not exhibit a significant discontinuity at 0.5. This is true both for cohabiting couples that eventually get married and for those who never marry. Once again, the only exception is couples formed at the workplace, for whom the drop is estimated to be 11.3%.

3.3 Separation and divorce

Figure 3 shows how the probability that a couple separates varies depending on the initial distribution of earnings within the household. Interestingly, the likelihood to separate exhibits a U-shape. The separation rate is the highest among couples with very unequal earnings, either in favor of men or in favor of women. The lowest separation rates are among couples where the woman earns about 40-45% of the total family earnings. There is no discontinuity in the probability of separation around the equal earnings point. This pattern is difficult to reconcile with the gender identity norm exhibiting a discontinuity at 0.5.⁸

⁸Newman and Olivetti (2018) propose a possible explanation. They argue that the increased bargaining flexibility in two-earner marriages makes them more durable than single-earner households.



Figure 3: Probability of couple separation, by the initial female share of household earnings

Notes: FLEED, 1988-2014. The sample includes couples that were initially identified as a couple based on their cohabitation status and observed for at least 15 years.

3.4 Who does exhibit the discontinuity and the spike at 0.5?

3.4.1 Co-working spouses

Figure 4 reports the relative earnings distribution separately for co-working and not-co-working couples using information for all years of the relationship. In the subsample of non-co-working spouses, we do not observe any discontinuity or missing mass of couples with women just outearning their husbands (panel a). As shown in Table 1, only 0.01% of these households are formed by equally earning spouses, and the estimate of the discontinuity at 0.5 is -0.2%, with the 95% confidence interval between 0.6% and -1.0%.

Instead, among co-working spouses, there is an excess mass of couples with partners having identical earnings and a sharp drop in the estimated limits of the density function at 0.5 (panel b). 6.2% of co-working spouses have the same earnings and the estimated drop at 0.5 is equal to 41%.

3.4.2 Self-employment versus working in the same firm

Figure 5 shows the relative earnings distributions separately for self-employed couples and couples working in the same firm. The two groups exhibit different patterns. Among self-employed couples, there is a large excess mass and a discontinuity at 0.5, while among spouses working in the same firm, the share of spouses with identical earnings is small but there is still a discontinuity at 0.5.

In Table 1 we report the share of equal earners in each group and the McCrary estimate for the discontinuity. 13.4% of self-employed couples have identical earnings,

Table 1: Equal earners and the discontinuity at 0.5 in the distribution of the female share of household earnings, by co-working status

	1	2	3	4
	Number of observations	% of all observations	% equal earners	Log Distance at 0.5
All employed couples	16 676 004	100	0.89	-0.113 (0.003)
Different firms	12 972 527	77.8	0.01	-0.002 (0.004)
Same industry	195 304	1.2	0.02	-0.039 (0.017)
Different industries	12 704 413	76.2	0.01	-0.002 (0.004)
Missing info on industry	72 810	0.4	0.02	-0.013 (0.029)
Same firm or both self-employed	2 354 062	13.4	6.2	-0.405 (0.004)
Both self-employed	1 045 301	5.5	13.4	-0.617 (0.005)
Same firm	1 308 761	7.9	0.4	-0.094 (0.007)
Same establishment	767 962	4.6	0.6	-0.160 (0.008)
Different establishments	500 110	3.0	0.03	-0.021 (0.012)
Missing establishment codes	40 689	0.2	0.5	-0.168 (0.031)
Missing employer code	1 349 415	8.2	0.2	-0.022 (0.009)
Same industry	25 979	0.2	0.7	-0.101 (0.046)
Different industries	1 039 822	6.3	0.1	-0.013 (0.010)
Missing info on industry	283 614	1.7	0.2	-0.064 (0.018)

Note: FLEED, 1988-2014. The sample includes couples, in which both partners are employed and have positive earnings. The group *missing employer code* includes employed spouses who are not self-employed and for whom there is no information about the identity of the employer. Industry is coded along 4-digit categories between 1988 and 1992 and along 5-digit categories between 1993 and 2014. Column 4 shows the log-distance at 0.5 between the left and right limits of the density function estimated after excluding observations at 0.5 using the McCrary (2008) procedure with default bins and bandwidths. Standard errors are in parentheses.

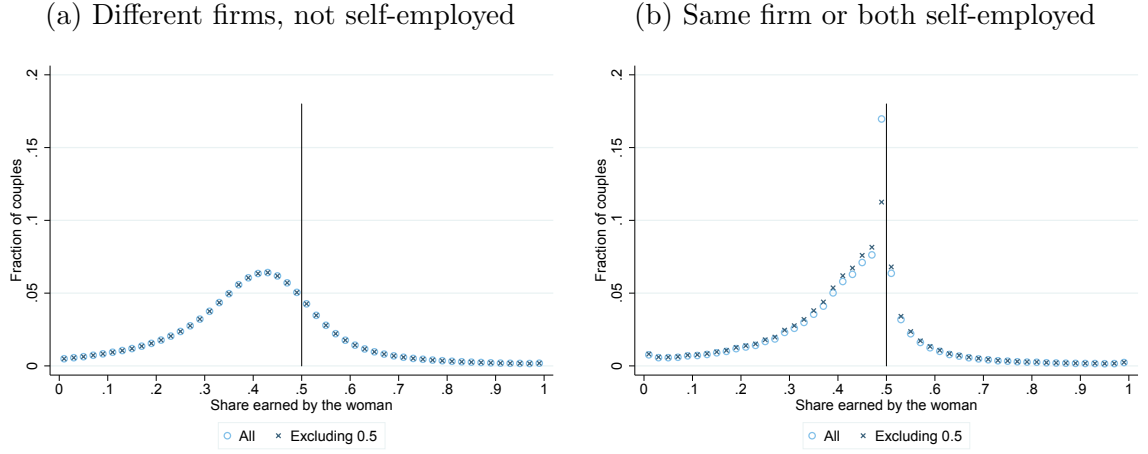


Figure 4: Relative earnings of women, by co-working status

Notes: FLEED, 1988-2014. The sample includes couples, in which both partners are employed and have positive earnings. Each dot is a fraction of couples in a 2% relative income bin; bins are right-closed. The dark-colored crosses show the fraction of couples in each bin after excluding households with identical earnings of both spouses.

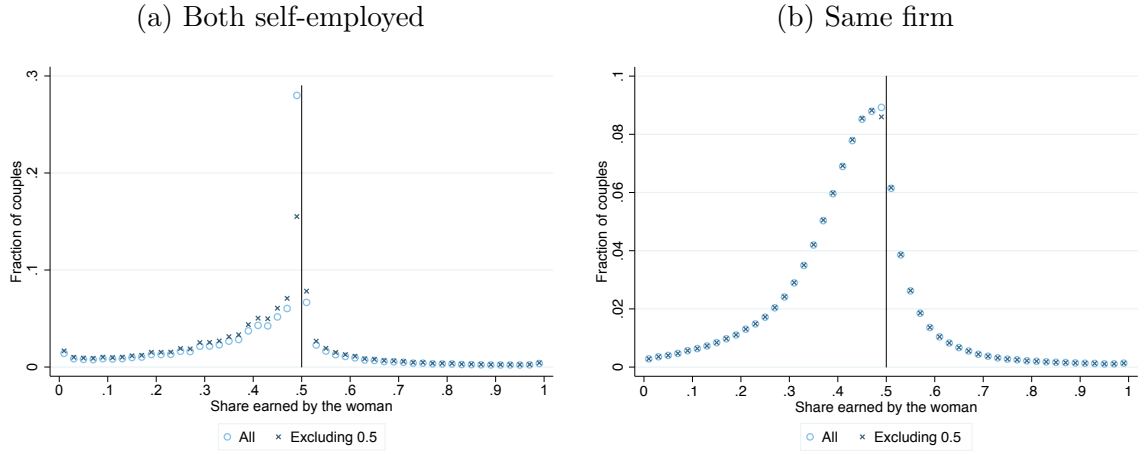


Figure 5: Relative earnings of women, co-working couples

Notes: FLEED, 1988-2014. The sample includes couples, in which both partners are employed and have positive earnings. Each dot is a fraction of couples in a 2% relative income bin; bins are right-closed. The dark-colored crosses show the fraction of couples in each bin after excluding households with identical earnings of both spouses.

and there are 62% more couples just below the 0.5 threshold than just above.⁹ In couples working in the same firm, 0.4% of spouses have identical earnings. The drop in the density is estimated to be 9% and it is largely driven by couples working in the same establishment, where the discontinuity is 16%. There is little evidence of discontinuity in households that work in different establishments of the same firm.

⁹Apart from couples with identical earnings, many self-employed couples have almost identical earnings. Nevertheless, there is a discrepancy between the left and the right limits of the density function at 0.5, even if observations in the vicinity of 0.5 are not taken into account for inferring those limits.

In Table 2 we investigate whether the above pattern is driven by couples working in a particular sector of the economy (agriculture, trade, manufacturing, education, or other sectors), in firms with a particular legal form (natural person, partnership, limited company, or other types), or in establishments of a certain size (less than 5 employees vs. 5-49 employees). In columns 1-3 we focus on self-employed couples. A high fraction of equal earners is present in all these different subgroups; however, equal earners are more prevalent in partnerships, reaching 35%. This is consistent with the legal default prescribing equal income-sharing in partnerships. The estimated discontinuity in the limits of the density function at 0.5 is statistically significant in all considered subsamples, but it is the largest in agriculture and small establishments.

Similarly, among spouses working in the same firm, the discontinuity is statistically significant in all sectors of the economy, across all legal forms and sizes of establishments (columns 4-6). The proportion of equal earners and the discontinuity are particularly large in agriculture, wholesale and retail trade, and in firms run as small family businesses. In none of these groups, the share of equal earners is larger than 3%.

3.4.3 Couples with different educational levels

According to the World Values Survey, less-educated individuals are more likely to agree with the statement that a woman should earn less than her husband to avoid problems. Bertrand, Kamenica and Pan (2015) employ spouses' educational level as a proxy for the prevalence of the gender norm. Consistently with the gender identity norm hypothesis, they find that the drop in the distribution of relative earnings is larger among less-educated couples.

In Table A2 and Figure A4 in the Online Appendix, we examine the interaction between spouses' educational level and co-working status. In Finland, as in the United States, the "missing" mass of couples with women just outearning their spouses is larger among less-educated couples. The estimated drop at 0.5 is 4.5% among college-educated couples, and it is 17% among less-educated ones. However, among both college and non-college educated couples, the discontinuity is present only among co-working spouses, who constitute about 12% of all higher educated couples and 17% of the less-educated couples. Among couples that do not work together, there is no excess mass or discontinuity at the 0.5 point, independent of the spouses' educational level. In this case, the estimate is a precise zero.

3.5 Evolution of the relative earnings in co-working couples

Next, we analyze the evolution of the distribution of relative earnings in co-working couples. First, we study the changes around the start of co-working in couples that initially were not working together. Second, we consider couples formed at the workplace and analyze the changes in the distribution around the start of cohabitation.

3.5.1 Couples who start co-working over time in the relationship

In Figure 6, we explore the evolution of the relative earnings distribution for spouses who start co-working after they formed a couple. The results are presented separately for couples who become jointly self-employed (plots on the left) and couples who start co-working in the same firm (plots on the right).

Table 2: **Equal earners and the discontinuity at 0.5 in co-working couples, by business characteristics**

	1	2	3	4	5	6
	Both self-employed			Same establishment		
	% of all observations	% equal earners	Log Distance at 0.5	% of all observations	% equal earners	Log Distance at 0.5
All	100	0.56	-0.160 (0.008)	100	13.4	-0.617 (0.005)
Sector:						
Agriculture	58.5	13.0	-0.710 (0.006)	1.9	1.8	-0.367 (0.040)
Wholesale, retail, repair of vehicles	5.6	13.0	-0.325 (0.025)	6.1	1.2	-0.233 (0.030)
Manufacturing	4.8	13.7	-0.471 (0.025)	32.7	0.2	-0.210 (0.014)
Education	1.0	12.7	-0.366 (0.047)	8.7	0.3	-0.150 (0.022)
Other	30.1	14.3	-0.408 (0.010)	50.6	0.7	-0.072 (0.009)
Legal form:						
Natural person	28.6	15.2	-0.543 (0.010)	1.2	2.2	-0.438 (0.064)
Partnership	2.3	35.1	-0.336 (0.032)	3.5	0.8	-0.420 (0.036)
Limited company	5.1	8.4	-0.278 (0.022)	65.7	0.6	-0.165 (0.010)
Government authority or enterprise	0.0	-	-	22.0	0.1	-0.154 (0.015)
Other	64.0	12.2	-0.682 (0.006)	7.5	0.7	-0.292 (0.023)
Establishment size:						
1-4 employees	59.6	11.7	-0.744 (0.007)	11.9	2.6	-0.217 (0.021)
5-49 employees	2.4	16.1	-0.389 (0.029)	25.2	0.7	-0.231 (0.016)
≥50 employees	0.0	-	-	40.1	0.05	-0.171 (0.012)
Missing info on size	37.9	16.0	-0.501 (0.007)	22.8	0.2	-0.146 (0.014)

Note: FLEED, 1988-2014. The sample includes couples, in which both partners are employed and have positive earnings. Column 3 and 6 show the log-distance at 0.5 between the left and right limits of the density function estimated after excluding observations at 0.5 using the McCrary (2008) procedure with default bins and bandwidths. Standard errors are in parentheses.

Let us first consider spouses that eventually become jointly self-employed. Consistent with our previous analysis, there is no discontinuity at the start of cohabitation (in blue on panel a). The year before the couple starts co-working, female earnings tend to be lower than at the beginning of the relationship (in red on panel a) but the shape of the relative earnings distribution is generally similar to the initial shape. The distribution of earnings changes radically when spouses become jointly self-employed (in red on panel c). A substantial fraction of women begin to have the same or very similar earnings as their husbands, while the proportion of all other combinations of spousal earnings decreases.¹⁰ Visually, there is no obvious discontinuity in the empirical density function if we ignore observations just around 0.5. Three years after spouses become self-employed, more couples start having relatively more similar earnings (in red on panel e), and the discontinuity at 0.5 between the left and right limits of the density function becomes salient. The McCrary estimate of this mismatch between the left and the right limit of the density function at 0.5 is 14%. In sum, the evidence for self-employed couples suggests that a sizable proportion of these couples tend to equalize earnings immediately after starting a family business. There is also a gradual convergence of earnings in these couples over time.

Let us now look at spouses who eventually start working in the same firm. At the start of cohabitation, the distribution of relative earnings in this sample is smooth around 0.5 (in blue on panel b). With time in the relationship, women start earnings relatively less, but the shape of the distribution of the female share is generally similar (in red on panel b). The distribution changes immediately after spouses start sharing the employer (panel d). There is a substantial increase in the share of couples where women earn between 30% and 55% of household earnings, and there is a thinning of both left and right tails of the distribution. In a small fraction of couples, women start earning exactly as much as their husbands, and a discontinuity appears at the 0.5 point. The estimated drop is equal to 15% and it is statistically significant. Overall, as a result of this convergence, there is an increase in the average relative earnings of women. The distribution remains stable in the following years (panel f).

This evidence is inconsistent with the discontinuity being caused by the gender identity norm, which predicts a decrease in the female share. Instead, we observe an increase both in the share of couples where men slightly outearn their wives and in the share of couples where women slightly outearn their husbands, but the increase in the former group is larger, and it generates a discontinuity around 0.5.

3.5.2 Couples formed at the workplace

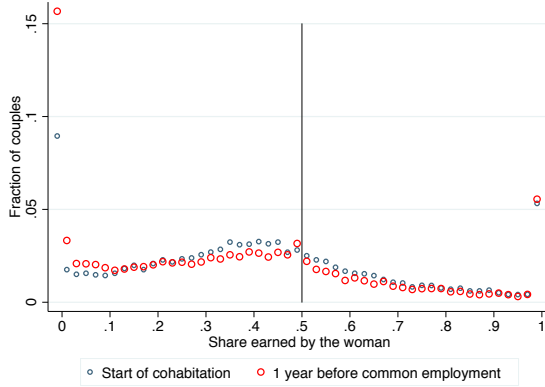
In Figure A5 in the Online Appendix, we analyze the evolution of the distribution of relative earnings in couples formed among individuals working in the same firm. We focus on the sample of couples who were already working together three years before the start of cohabitation and still work together one year after.

The distribution of the female share becomes substantially compressed around 0.5 from both sides of the distribution about two years before the start of cohabitation. The McCrary test detects a discontinuity at the point of equal earnings of about 16%. However, the distribution remains almost unchanged one year before and one year

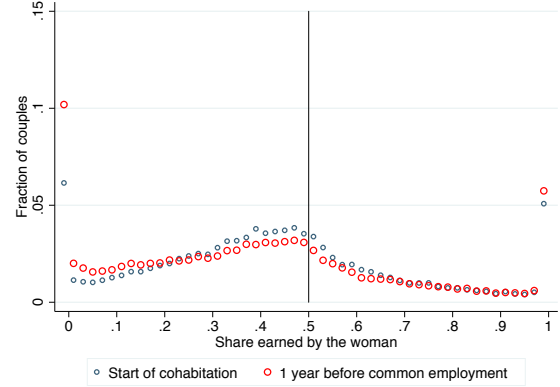
¹⁰Figure A3 in the Online Appendix compares the distribution of the initial female share in couples that eventually equalize their earnings with the corresponding distribution in all couples. It is clear that couples who equalize their earnings come from both sides of the distribution.

Panel A. The start of cohabitation and a year before the start of co-working

(a) Self-employed




(b) Same firm

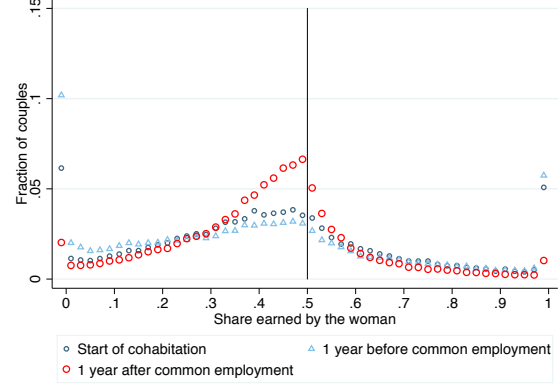


Panel B. A year after the start of co-working

(c) Self-employed




(d) Same firm



Panel C. 3 years after the start of co-working

(e) Self-employed



(f) Same firm




Figure 6: Dynamics of the relative earnings of women

Notes: FLEED, 1988-2014. The sample includes couples that start co-working during cohabitation, either as a self-employed couple or as wage earners in the same firm; it is restricted to couples that were initially identified as a couple based on their cohabitation status and are observed for at least 15 years. Each dot is a fraction of couples in a 2% relative income bin; bins are right-closed.

after co-working individuals start to cohabit. Overall, the dynamics are consistent with earnings convergence in co-working couples, either due to earnings compression at the firm level or due to information-sharing among future spouses already before the start of cohabitation.

3.6 Actual and potential earnings

If co-working “activates” the gender identity norm, women who start co-working with their spouses would earn, on average, less than comparable non-co-working women. In this section, we compare the evolution of earnings for women who started co-working with their husbands with the evolution of earnings for women who never worked together with their spouses, taking into account observable differences in predetermined characteristics. We also conduct the same analysis for men. To predict the counterfactual earnings of co-working spouses, we first estimate the following set of equations on the subsamples of non-co-working women and men:

$$Y_{i,k,t} = \beta_0^k + \mathbf{X}_i\beta_1^k + \mathbf{D}_t\beta_2^k + \epsilon_{i,k,t} \quad (1)$$

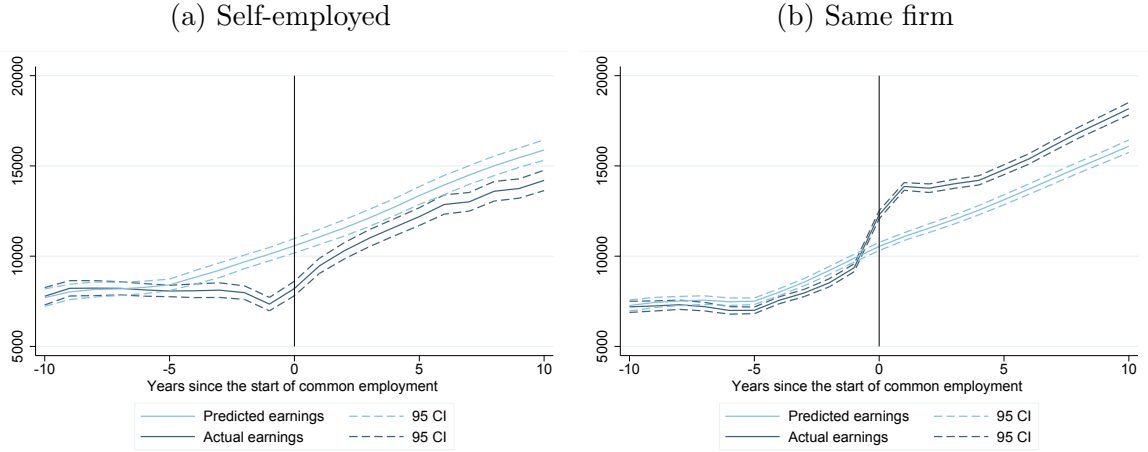
where $Y_{i,k,t}$ represents real earnings in year t of individual i who has cohabited with her or his spouse for k years. X_i is a vector of predetermined characteristics measured the year before the start of cohabitation, which includes the main activity (employed, unemployed, student, pensioner, conscript, unemployment pension, other inactive), industry (20 categories), occupation (10 categories), establishment size (10 categories), earnings, indicator for zero earnings, age dummies, nationality (Finnish or other), family structure (7 categories), interaction between educational level (6 categories) and field (8 categories), and region of residence. \mathbf{D}_t are year dummies.

We then use the estimates from this model for out-of-the-sample prediction of counterfactual earnings for individuals who choose to co-work with their spouses at some point in the relationship. In panels (a) and (c) of Figure 7, we compare the earnings of co-working spouses with their counterfactual using a ten-year window around the start of joint self-employment. The earnings of men and women who at some point in their relationship become jointly self-employed with their spouses are initially similar to the earnings of other individuals with similar characteristics. Before becoming jointly self-employed, both men and women tend to experience a negative earnings shock, but the shock is relatively larger in the case of women.¹¹ The trajectory of men and women changes after the start of self-employment. The earnings of men drop substantially the year when they become self-employed. Instead, the earnings of women increase and almost reach the counterfactual level. As we show in Figure A6 in the Online Appendix, this catch-up in female earnings is mostly due to an increase in employment rates.

The fact that family earnings remain below the counterfactual level after the start of joint self-employment is consistent with couples choosing self-employment as a way to recover from a negative employment shock. Some families with children may be also willing to accept lower earnings in exchange for the flexibility of work schedules in self-employment. There is also evidence suggesting that self-employed couples shift income between wages and dividends for fiscal reasons (Harju and Matikka, 2016).

¹¹As we show in Figure A6 in the Online Appendix, two years before the start of co-working, the employment rate of women falls below their potential. The aggregate earnings differential is both due to lower employment and lower earnings conditional on employment.

A. Women



B. Men

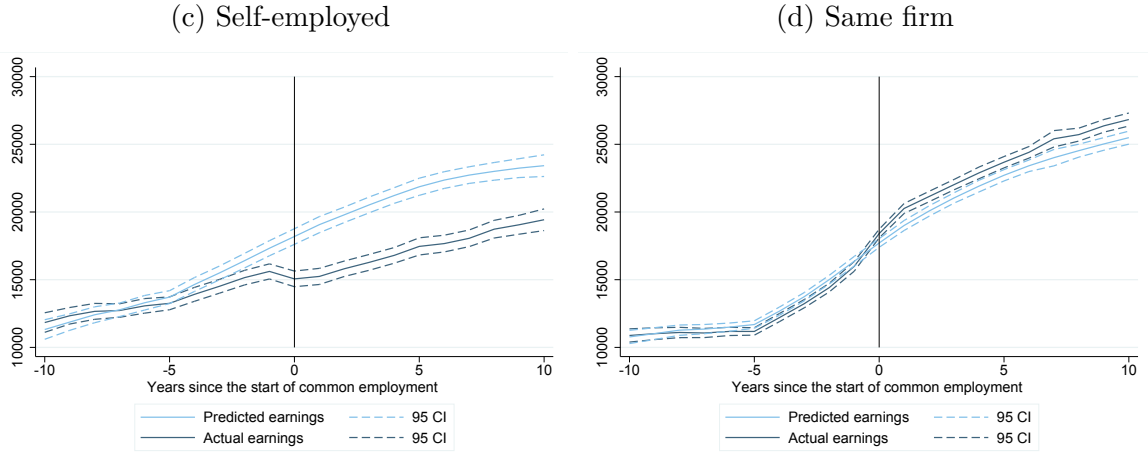


Figure 7: Actual and predicted earnings in co-working couples

Notes: FLEED, 1988-2014. Predicted earnings are obtained using out-of-the-sample prediction from a model estimated on a sample of individuals who never worked together with their spouses by regressing individual earnings in each particular year into the relationship on predetermined main activity, industry, occupation, earnings, indicator for zero earnings, age dummies, nationality, family structure, education level and field, and region of residence as observed the year before the start of cohabitation. Dashed lines indicate 95% confidence intervals.

Panels (b) and (d) of Figure 7 show the analysis, respectively, for women and men who at some point in the relationship start working in the same firm. While initially, the earnings of individuals in future co-working couples are similar to the counterfactual level, before the start of co-working, female earnings fall slightly short of the counterfactual level.¹² After the start of co-working, women start earning 25% more than non-co-working women with similar predetermined characteristics, while men's earnings become about 7% higher than the counterfactual level. As we show in Figure A7 in the Online Appendix, the increase in male earnings can be entirely attributed to the increase in employment probability, while female earnings increase both due to

¹²As we show in Figure A7 in the Online Appendix, the disadvantage in earnings among co-working women is due to a slightly lower employment rate.

the elevated probability of employment and earnings conditional on employment.

The increase in female earnings above the potential predicted by predetermined observable characteristics is consistent with several hypotheses. Women in co-working couples may have a stronger bargaining position in wage negotiations or they may be favored. Couples may also choose to co-work only if co-working implies a substantial improvement in female earnings.

3.7 Pairs of unrelated co-working women and men

We analyze the evolution of relative earnings in placebo couples, which we construct by randomly matching unrelated co-working women and men. Specifically, we first selected a 5% random sample of all possible pairs of co-working men and women who were never observed as cohabiting or married in our data. We then imposed two restrictions on the educational level and age differences between individuals in these pairs so that they correspond to typical matches observed in the marriage market. First, we considered pairs formed by individuals who either both have achieved, at most, high school education or both have at least some college education. Second, we limited the analysis to pairs in which men are between four and zero years older than women. Over two-thirds of actual couples in Finland satisfy the first condition and over half the second. After applying the age restriction, the average age gap in fictitious couples is two years, the same as in actual ones. Finally, to be able to study the dynamics of relative earnings over time, we considered individuals who co-worked for at least 15 years in the same plant. About 80% of these placebo couples work in firms with at least one cohabiting couple.

Figure 8 shows the evolution of the distribution of the female share in the placebo couples. One year before individuals start to co-work in the same firm, the distribution of the female share of earnings is smooth over 0.5. One year after the start of co-working, the distribution becomes substantially more compressed, but no discontinuity at 0.5 is detected (see estimates in Table A3 in the Online Appendix). Over time, the earnings of individuals continue converging, and a cliff at 0.5 emerges. After five years of co-working, the estimated drop at 0.5 is 14.6%, significant at 1%. A similar drop is still observed after 15 years.

The existence of a discontinuity among placebo co-working couples suggests that the observed discontinuity among co-working spouses is, at least partly, due to firm-level earnings compression.

3.8 Discontinuity at other points of the distribution

We explore whether there exist bunching and discontinuity at other points of the relative earnings distribution, besides point 0.5. In particular, we estimate discontinuities to the right of each 2-percentage point bin between 0.1 and 0.9. To facilitate the comparison, we keep observations at 0.5. We consider separately the following three groups: non-co-working couples, self-employed couples, and couples working in the same firm. In the case of non-co-working couples, we do not expect any discontinuities at any point of the distribution. For couples working in the same firm, we only expect a discontinuity at the 0.5 point. Finally, in self-employed couples, the existence of ad-hoc rules for entrepreneurial income sharing may generate spikes and discontinuities at other salient bin-separating points.



Figure 8: Evolution of the relative earnings in pairs of randomly matched men and women co-working in the same plant

Notes: FLEED, 1988-2014. The sample consists of pairs formed by randomly matched unrelated men and women co-working in the same plant. The sample is restricted to pairs in which individuals have a similar educational level (high school or lower versus some college or higher), in which men are between zero and four years older than women, and who co-worked in the same firm for at least 15 years. Each dot is a fraction of pairs in a 2% relative income bin; bins are right-closed.

We report our results in Figure A8 in the Online Appendix. In the sample of non-co-working couples, out of 40 estimates, only one is significant at 5% level (p-value=0.020). In the sample of couples working in the same firm, a large discontinuity at 0.5 affects the estimates of the limits of the density function in other neighboring points, however, there is only one estimate outside the vicinity of 0.5 significant at the 5% (p-value=0.016). Given the number of tested parameters, these discontinuities are likely to be false positives and, overall, we cannot reject that the density function is smooth away from 0.5. For self-employed couples, we observe significant bunching of observations in the vicinity of several points, most importantly, around 0.4, 0.3, 0.33, 0.2, 0.25, 0.1, 0.6, 0.7 and 0.8. This bunching pattern is consistent with couples using simple ad-hoc rules for entrepreneurial income sharing. All points with major bunching are detected by the McCrary test. Again, bunching affects the estimates in other neighboring points. However, the disturbance created by bunching at 0.5 is particularly large.

3.9 Evidence from the United States

The above evidence suggests that, in Finland, the discontinuity and the excess mass at 0.5 are due to the equalization of earnings in self-employed couples and to convergence of earnings in couples co-working in the same firm. In this section, we discuss whether this explanation may also apply to the United States.

Unfortunately, the SIPP/SSA/IRS dataset used by Bertrand, Kamenica and Pan

(2015) does not include information on the firm where individuals work or indicate whether they are self-employed. Information from other sources suggests that, in the United States, 11-13% of wage-earning spouses work for the same employer and about 3% of working couples are simultaneously self-employed.¹³

The phenomenon of earnings equalization is less relevant in the United States than in Finland. In the United States, only 0.3% of couples have identical earnings, compared to 0.9% in Finland. Some American self-employed couples may want to equalize earnings to simplify accounting or, perhaps, to avoid within-family negotiations. However, in the United States, unlike in Finland, there are no legal defaults for income sharing in partnerships, and households can jointly file their income tax declarations.

For couples co-working in the same firm, the impact of earnings compression is likely to have a similar effect as in Finland. To assess the relevance of income convergence in co-working couples, we use the SIPP/SSA/IRS dataset and we proxy whether spouses work together using available information on industry and occupation. It seems reasonable to expect that the share of co-working couples is substantially higher among couples working in the same industry and occupation.¹⁴ Instead, couples working in different industries are unlikely to work in the same firm; although some self-employed couples may be included in this group.¹⁵

We observe that around 20% of all couples work in the same industry and occupation, while 60% of couples work in different industries. Figure 9 shows the distribution of relative earnings separately for these two groups of couples. The drop in the distribution at 0.5 is significantly larger among couples working in the same industry and occupation. According to the McCrary test, the estimate of the drop is 14%: about twice as large as the drop observed in the overall population. This evidence suggests that factors leading to earnings convergence in co-working couples are also likely to play a significant role in explaining the existence of a discontinuity in the United States.

4 Conclusions

We study the underlying causes for the existence of a sharp drop to the right of 0.5 in the distribution of households according to the share of total earnings earned by the wife. This discontinuity, which was originally observed by Bertrand, Kamenica and Pan (2015) among U.S. households, has been attributed to the existence of a gender identity norm prescribing a breadwinner role to men. According to this hypothesis, couples where women even slightly outearn men are significantly less likely to be formed, more likely to divorce, and women in this couples tend to reduce their labor supply to avoid outearning their spouses.

We propose an alternative explanation. We argue that the discontinuity to the right

¹³Hyatt (2015) estimates the share of co-working couples among wage earners using data from the U.S. Census Bureau’s Longitudinal Employer-Household Dynamics (LEHD) program. The estimate of the proportion of simultaneously self-employed couples is our own based on data from the American Community Survey (ACS) 2008-2011.

¹⁴According to Hyatt (2015), among couples working in the same narrowly defined Census industry and Census occupation, the proportion of co-working couples is 83%. The level of disaggregation available in the SIPP/SSA/IRS dataset is much lower than in the study by Hyatt. Industries are classified into only four categories, and occupations are divided into three categories. Information on industry and occupation is missing for 15% of couples.

¹⁵If we apply SIPP/SSA/IRS industry classification to the ACS data, we observe that 24% of jointly self-employed couples report to work in different industries.

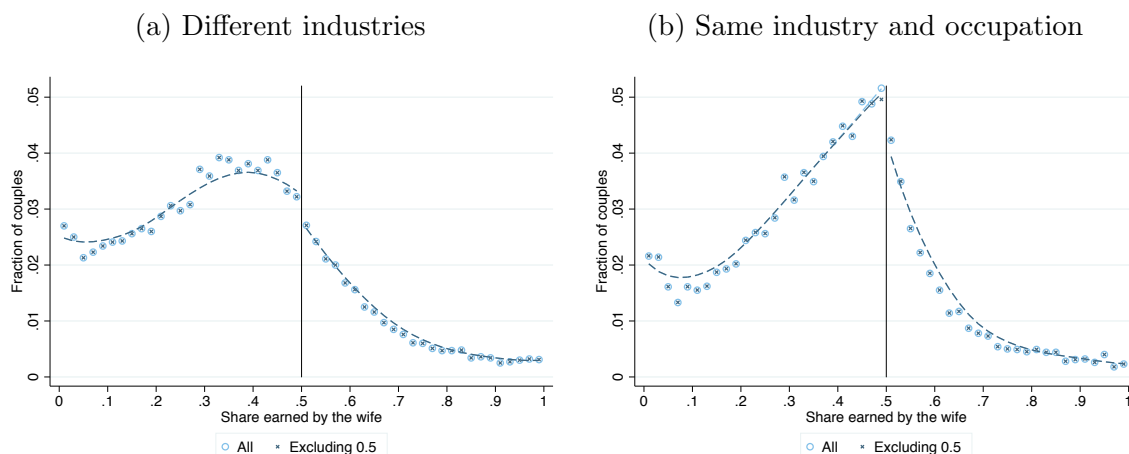


Figure 9: Relative earnings of women, U.S. households

Notes: SIPP/SSA/IRS Completed Gold Standard Files, 1990-2004. The sample includes married couples with both partners aged between 18 and 65 years and receiving positive earned income. Each dot indicates the fraction of couples in a 2% relative income bin; bins are right-closed. The dashed lines are the lowess smoothers applied to the distribution allowing for a break at 0.5.

of 0.5 can emerge if some couples tend toward earnings equalization or convergence. To test this hypothesis, we exploit the rich employer-employee linked data from Finland. We find overwhelming support in favor of the idea that the discontinuity is caused by earnings equalization in self-employed couples and earnings convergence among spouses working together. We show that the discontinuity is not generated by selective couple formation or separation and it arises only among self-employed and co-working couples who account for 15% of the population.

Self-employed couples are responsible for most observations with spouses reporting identical earnings. When couples start being self-employed, both sides of the distribution tend to equalize earnings, perhaps because earnings equalization helps couples to reduce income tax payments, facilitate accounting, or avoid unnecessary within-family negotiations. Large spikes emerge not only at 0.5 but also at other round shares signaling the prevalence of ad-hoc rules for entrepreneurial income sharing in couples. Self-employment is associated with a fall of household earnings below the level predicted by individuals' predetermined characteristics, but this drop is mainly due to a decrease in male earnings, with women being relatively better off.

In the case of couples who work together in the same firm, there is a compression of the earnings distribution toward 0.5 both on the right and on the left of 0.5. As a result, there is an increase both in the share of couples where men slightly outearn their wives and in the share of couples where women slightly outearn their husbands. Since the former group is larger, earnings compression leads to a detection of a discontinuity. Notably, we observe a similar earnings compression and a discontinuity in relative earnings among fictitious couples that we construct by randomly matching unrelated women and men co-working in the same firm. This 'placebo' suggests that firm-level factors contribute to the emergence of the discontinuity. We also observe that, in couples who start co-working, household earnings tend to increase above the level predicted by spouses' observable characteristics and this increase is mainly driven by a jump in female earnings above their earnings potential. This increase in female earnings

is consistent with women in co-working couples gaining from joint negotiations with the employer, or with couples choosing to co-work only if it implies a substantial improvement in female earnings. Overall, the evidence suggests that the observed discontinuity is not due to the existence of a social norm that limits the income of married women. Paradoxically, it is a result of a phenomenon, co-working couples, that helps women to have higher earnings.

Due to data limitations, we are not able to provide comparable evidence for the United States, but we find that the discontinuity is twice as large among households with spouses working in the same industry and occupation, and hence having a higher likelihood of being co-employed. Arguably, this pattern also supports the relevance of the hypothesis of earnings convergence in households with co-working spouses as an explanation for the discontinuity in the United States. Nonetheless, a more comprehensive analysis using administrative data on individual earnings and employment histories needs to be conducted to confirm the validity and significance of this hypothesis.

While our results suggest that the discontinuity should not be considered as evidence for the existence of the gender identity norm, we would like to emphasize that this does not imply that the norm does not play an important role in the marriage market and in women’s labor supply decisions. It is possible that the norm only gradually gains importance with the increase in the relative earnings of women, and there is no sharp discontinuity or kink in the utility function immediately to the right of the point with equal earnings of spouses. Further research is needed to test this hypothesis.

References

- Bertrand, Marianne, Emir Kamenica, and Jessica Pan.** 2015. “Gender Identity Norms and Relative Income within Households.” *Quarterly Journal of Economics*, 130(2): 571–614.
- Binder, Ariel J., and David Lam.** 2018. “Is There a Male Breadwinner Norm? The Hazards of Inferring Preferences From Marriage Market Outcomes.” *NBER working paper 24907*.
- Breza, Emily, Supreet Kaur, and Yogita Shamdasani.** 2018. “The Morale Effects of Pay Inequality.” *Quarterly Journal of Economics*, 133(2): 611–663.
- Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard.** 2019. “Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior.” *American Economic Review*, 109(2): 620–663.
- Harju, Jarkko, and Tuomas Matikka.** 2016. “Business owners and income-shifting: evidence from Finland.” *Small Business Economics*, 46(1): 115–136.
- Hederos, Karin, and Anders Stenberg.** 2015. “Gender Identity and Relative Income within Households: Evidence from Sweden.” *IZA Discussion paper 9533*.
- Hyatt, Henry R.** 2015. “Co-Working Couples and the Similar Jobs of Dual-Earner Households.” *US Census Bureau Center for Economic Studies Paper CES-WP-15-23*.

- Inglehart, Ronald, Christian Haerpfer, Alejandro Moreno, Christian Welzel, Kseniya Kizilova, Jaime Diez-Medrano, Marta Lagos, Pippa Norris, Eduard Ponarin, Bi Puranen, and others. (eds).** 2014. “World Values Survey: Round Three - Country-Pooled Datafile 1995-1998.” Version: <http://www.worldvaluessurvey.org/WVSDocumentationWV3.jsp> (accessed August 25, 2020). Madrid: JD Systems Institute.
- Kabatek, Jan, Arthur van Soest, and Elena Stancanelli.** 2014. “Income taxation, labour supply and housework: A discrete choice model for French couples.” *Labour Economics*, 27: 30–43.
- Kleven, Henrik, Claus Kreiner, and Emmanuel Saez.** 2009. “The optimal income taxation of couples.” *Econometrica*, 72(2): 537–560.
- LaLumia, Sara.** 2008. “The effects of joint taxation of married couples on labor supply and non-wage income.” *Journal of Public Economics*, 92: 1698–1719.
- McCrary, Justin.** 2008. “Testing for manipulation of the running variable in the regression discontinuity design: a density test.” *Journal of Econometrics*, 142: 698–714.
- Murray-Close, Marta, and Misty L. Heggeness.** 2018. “Manning up and womaning down: How husbands and wives report their earnings when she earns more.” *SESHD Working Paper 2018-20*.
- Newman, Andrew F., and Claudia Olivetti.** 2018. “Career Women and the Durability of Marriage.” *Boston University, mimeo*.
- Saez, Emmanuel, Benjamin Schoefer, and David Seim.** 2019. “Payroll Taxes, Firm Behavior, and Rent Sharing: Evidence from a Young Workers’ Tax Cut in Sweden.” *American Economic Review*, 109(5): 1717–1763.
- Schuetze, Herbert J.** 2006. “Income splitting among the self-employed.” *The Canadian Journal of Economics*, 39(4): 1195–1220.
- Statistics Finland.** 2019. “Finnish Longitudinal Employer-Employee Data (FLEED).” Helsinki: Statistics Finland.
- Stephens, Melvin Jr, and Jennifer Ward-Batts.** 2004. “The impact of separate taxation on the intra-household allocation of assets: evidence from the UK.” *Journal of Public Economics*, 88: 1989–2007.
- U.S. Census Bureau.** 2015*a*. “SIPP Completed Gold Standard Files. Version 6.02.” United States Department of Commerce, Washington, D.C.
- U.S. Census Bureau.** 2015*b*. “SIPP Synthetic Beta. Version 6.02.” United States Department of Commerce, Washington D.C.; Cornell University, Synthetic Data Server, Ithaca, NY.

A Online Appendix

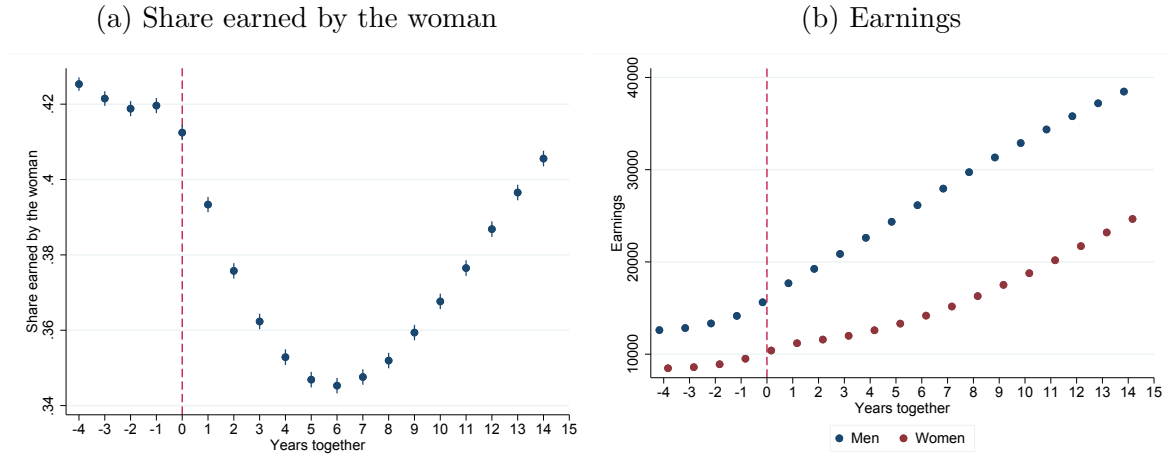


Figure A1: Evolution of household earnings around the start of cohabitation
Notes: FLEED, 1988-2014. The sample includes couples that were initially identified as a couple based on their cohabitation status and observed for at least 15 years.

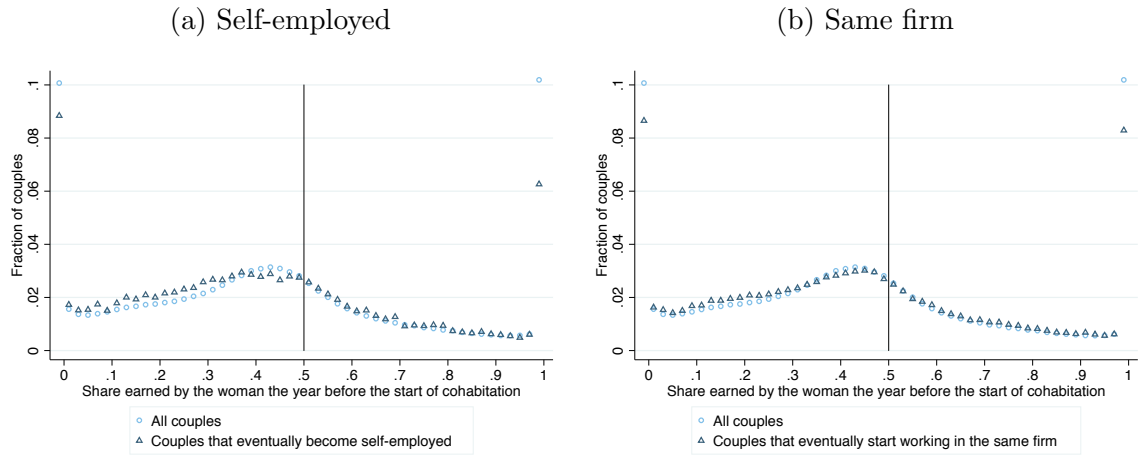


Figure A2: Self-selection into co-working
Notes: FLEED, 1988-2014. The sample includes couples with both partners being employed and receiving positive earned income at the year of marriage. Each dot is a fraction of couples in a 2% relative income bin; bins are right-closed.

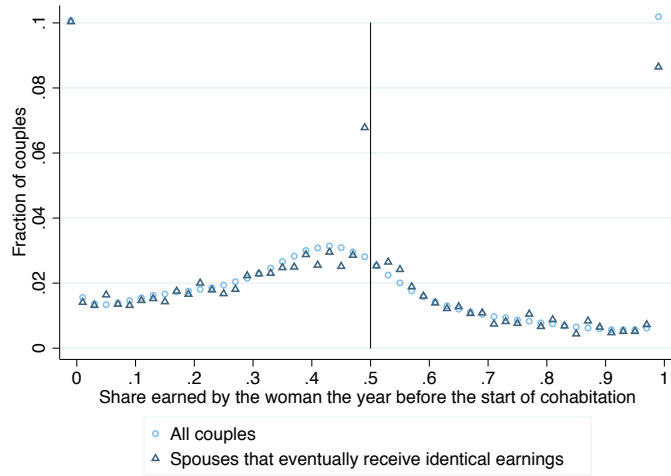


Figure A3: Same earnings and the initial female share

Notes: FLEED, 1988-2014. The sample includes couples with both partners being employed and receiving positive earned income at the year of marriage. Each dot is a fraction of couples in a 2% relative income bin; bins are right-closed.

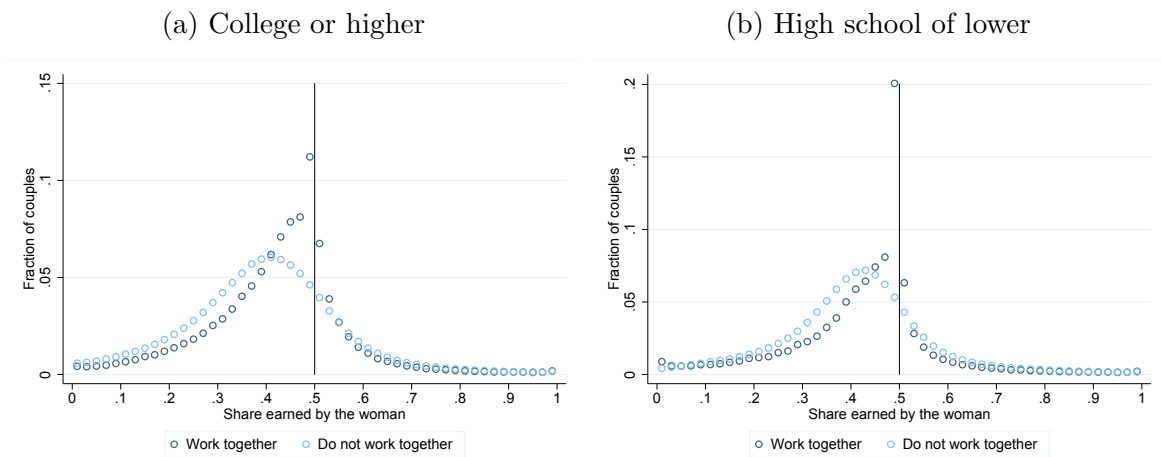


Figure A4: Relative earnings of women, by educational level and co-working status

Notes: FLEED, 1988-2014. The sample includes couples with both partners being employed and receiving positive earned income at the year of marriage. Each dot is a fraction of couples in a 2% relative income bin; bins are right-closed.

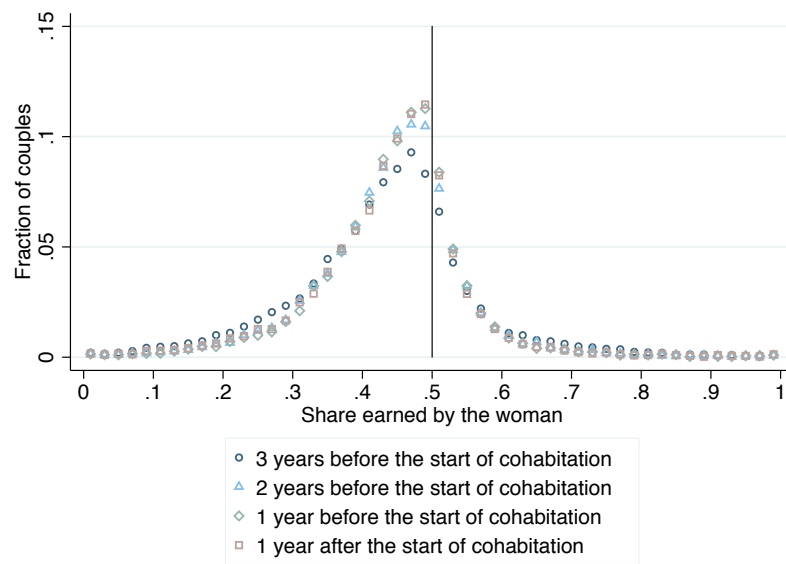
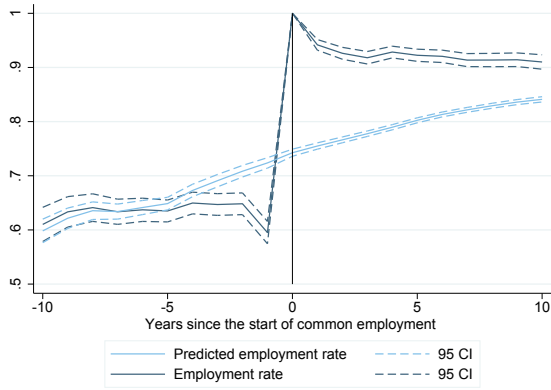


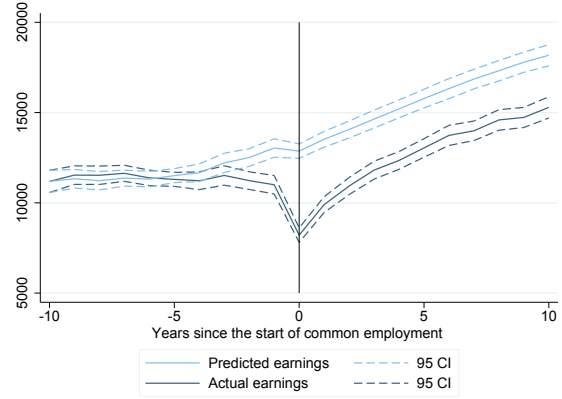
Figure A5: Evolution of the relative earnings in couples formed in the workplace
Notes: FLEED, 1988-2014. The sample includes couples that started to cohabit after being coworkers in the same firm for at least 3 years.

A. Women

(a) Employment rate




(b) Earnings conditional on being employed



B. Men

(c) Employment rate



(d) Earnings conditional on being employed

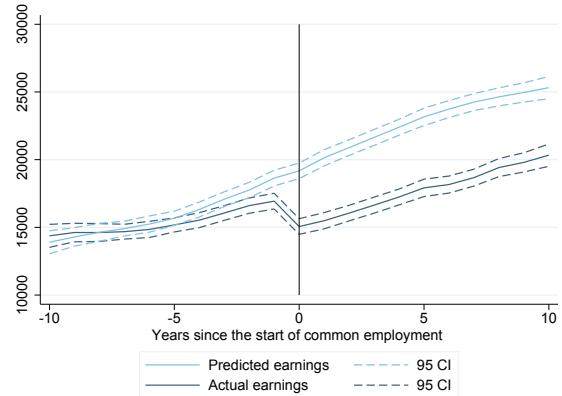


Figure A6: Actual and predicted employment and wages, self-employed couples
Notes: FLEED, 1988-2014. Predicted earnings are obtained using out-of-the-sample prediction from a model estimated on a sample of individuals who never worked together with their spouses by regressing individual labor market outcomes in each particular year into the relationship on predetermined main activity, industry, occupation, earnings, indicator for zero earnings, age dummies, nationality, family structure, education level and field, and region of residence as observed the year before the start of cohabitation. Dashed lines indicate 95% confidence intervals.

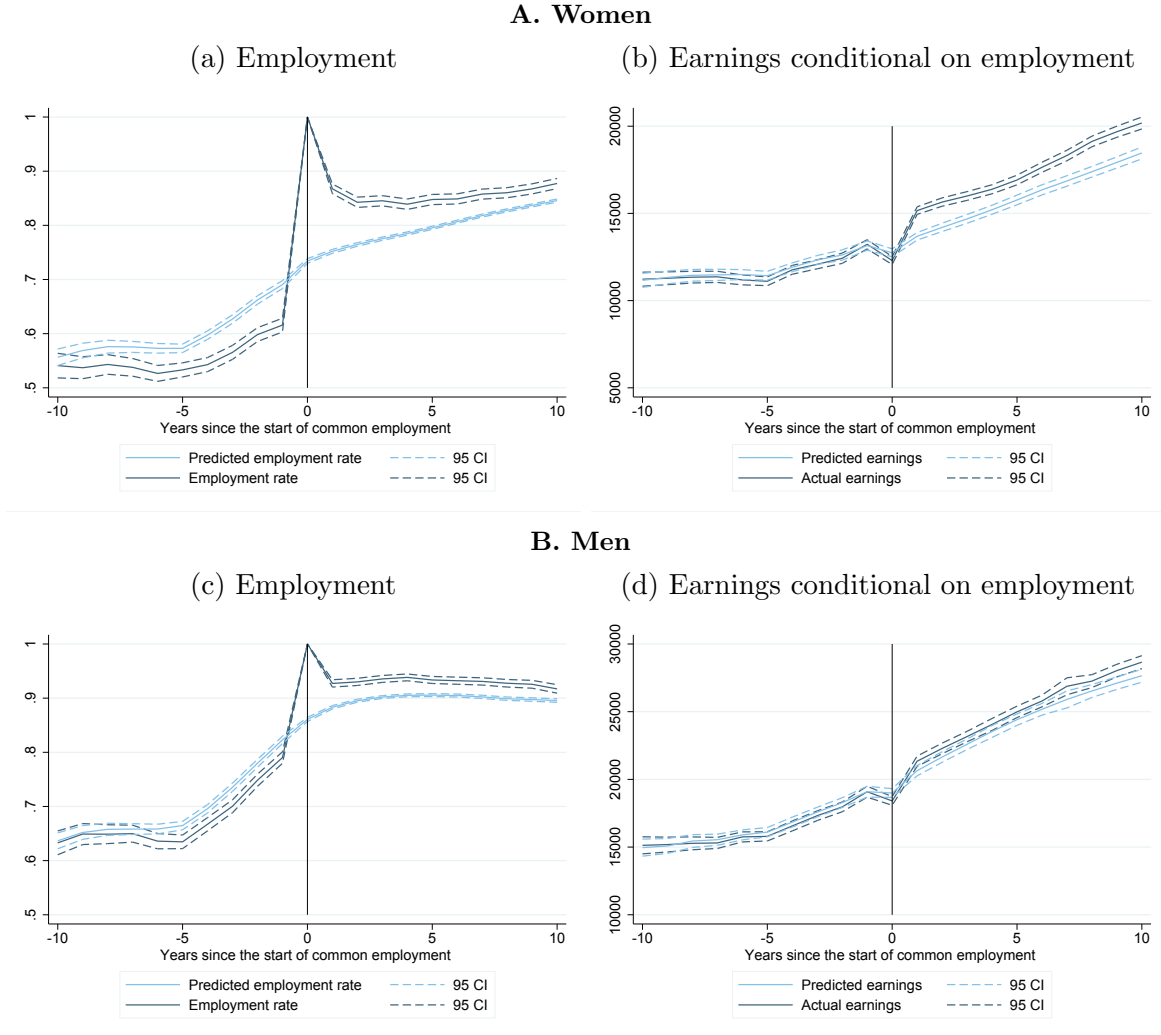



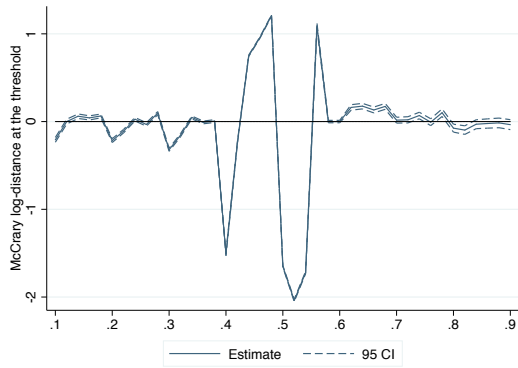
Figure A7: Actual and predicted employment and wages, spouses co-working in the same firm

Notes: FLEED, 1988-2014. Predicted earnings are obtained using out-of-the-sample prediction from a model estimated on a sample of individuals who never worked together with their spouses by regressing individual labor market outcomes in each particular year into the relationship on predetermined main activity, industry, occupation, earnings, indicator for zero earnings, age dummies, nationality, family structure, education level and field, and region of residence as observed the year before the start of cohabitation. Dashed lines indicate 95% confidence intervals.

(a) Not co-working



(b) Both self-employed



(c) Same firm




Figure A8: McCrary estimates in other bin-separating points

Notes: McCrary (2008) estimates of the discontinuity in the distribution of female relative earnings to the right of each 2-percentage point bin between 0.1 and 0.9, with 95% confidence intervals (CI). All estimates are based on default bandwidths and bin-sizes.

Table A1: **McCrary tests, marriage versus start of cohabitation**

	1	2	3	4
	Number of observations	% of observations	% equal earners	Log Distance at 0.5
First year of marriage	415 666	100	0.2	-0.041 (0.015)
Not co-working couples	342 929	82.5	0.00	-0.029 (0.017)
Co-working couples	44 809	10.8	1.9	-0.193 (0.027)
First year of cohabitation	654 087	100	0.04	-0.014 (0.012)
Couples that eventually marry	233 322	35.7	0.04	-0.016 (0.019)
Couples that never marry	420 589	64.3	0.04	-0.023 (0.014)
Not co-working couples	552 301	84.4	0.00	-0.007 (0.013)
Co-working couples	64 564	9.9	0.3	-0.113 (0.024)

Note: FLEED, 1988-2014. Households in which partners have identical earnings are excluded. Column 4 shows the log-distance at 0.5 between the left and right limits of the density function estimated after excluding observations at 0.5 using the McCrary (2008) procedure with default bins and bandwidths. Standard errors are in parentheses.

Table A2: McCrary tests, by educational level and co-working status

	1	2	3	4
	Number of observations	% of observations	% of equal earners	Log Distance at 0.5
Both college or higher	3 519 220	100	0.2	-0.045 (0.006)
Different firms	2 914 072	82.8	0.01	-0.003 (0.008)
Both self-employed	75 176	2.1	9.5	-0.449 (0.018)
Same firm	343 112	9.7	0.2	-0.094 (0.012)
Missing employer code	186 860	5.3	0.1	-0.071 (0.021)
Both high school or lower	8 524 389	100	1.3	-0.171 (0.003)
Different firms	6 279 087	73.7	0.01	-0.002 (0.005)
Both self-employed	787 423	9.2	14.0	-0.700 (0.006)
Same firm	652 832	7.7	0.4	-0.151 (0.009)
Missing employer code	805 047	9.4	0.2	-0.039 (0.011)

Note: FLEED, 1988-2014. The sample includes couples with both partners being employed and receiving positive earnings. Households in which partners have identical earnings are excluded. The group *missing employer code* includes employed spouses who are not self-employed and for whom there is no information about the identity of the employer. Column 4 shows the log-distance at 0.5 between the left and right limits of the density function estimated after excluding observations at 0.5 using the McCrary (2008) procedure with default bins and bandwidths. Standard errors are in parentheses.

Table A3: McCrary tests, randomly matched unrelated women and men co-working in the same plant

	% of equal earners	Log Distance at 0.5
All pairs	0.02	-0.055 (0.014)
Only pairs that eventually are observed co-working for at least 15 years:		
1 year before the start of co-working	0.01	-0.035 (0.060)
1 year of co-working	0.01	-0.052 (0.038)
5 years of co-working	0.04	-0.146 (0.034)
10 years of co-working	0.06	-0.054 (0.023)
15 years of co-working	0.03	-0.129 (0.030)

Note: FLEED, 1988-2014. The tables shows the log-distance at 0.5 between the left and right limits of the density function estimated after excluding observations at 0.5 using the McCrary (2008) procedure with default bins and bandwidths. Standard errors are in parentheses. The sample consists of pairs formed by randomly matched unrelated men and women co-working in the same plant. The sample is restricted to pairs in which individuals have a similar educational level (high school or lower versus some college or higher) and in which men are between zero and four years older than women.