

Total Work, Gender and Social Norms in EU and US Time Use

Michael C Burda, Daniel S Hamermesh, Philippe Weil

► **To cite this version:**

Michael C Burda, Daniel S Hamermesh, Philippe Weil. Total Work, Gender and Social Norms in EU and US Time Use. Burda, Michael C, Boeri, Tito, Kramarz, Francis Burda, Michael C. Working Hours and Job Sharing in the EU and USA: Are Europeans Lazy? Or Americans Crazy?, Oxford University Press, pp.11-91, 2008. hal-00972821

HAL Id: hal-00972821

<https://hal-sciencespo.archives-ouvertes.fr/hal-00972821>

Submitted on 3 Apr 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

November 2008

TOTAL WORK, GENDER AND SOCIAL NORMS

Michael Burda, Daniel S. Hamermesh and Philippe Weil*

*Professor of economics Humboldt-Universität zu Berlin, research associate, CEPR and IZA; Sue Killam Professor of economics, University of Texas at Austin, research associate, IZA and NBER; professor of economics, Université Libre de Bruxelles and Institut d'Études Politiques de Paris, research associate, CEPR and NBER. We thank Olivier Blanchard, Tito Boeri, Micael Castanheira, Deborah Cobb-Clark, Raquel Fernandez, Georg Kirchsteiger, Christopher Pissarides, and participants in seminars at several universities and conferences for helpful comments, and Rick Evans and Juliane Scheffel for their expert research assistance.

JEL Codes: J22; J16; D13

Keywords: time use; gender differences; household production; paid work

ABSTRACT

Using time-diary data from 27 countries, we demonstrate a negative relationship between real GDP per capita and the female-male difference in total work time—the sum of work for pay and work at home. We also show that in rich non-Catholic countries on four continents men and women do the same amount of total work on average. Our survey results demonstrate that labor economists, macroeconomists, sociologists and the general public consistently believe that women perform more total work. The facts do not arise from gender differences in the price of time nor from differences in intra-family bargaining: Gender equality is not associated with marital status, and most of the variance in gender total work differences arises from within-couple differences. A theory of social norms could account for within-education group and within-region gender differences being smaller than inter-group differences. It is consistent with cross-national evidence from the World Values Surveys and various sets of microeconomic data.

1 Introduction

It is well-known that men engage in more market work—have higher participation rates and longer workweeks conditional on participation—than women. What has not been thoroughly examined, and indeed has been almost ignored by economists, is the issue of gender differences in the total amount of work—the sum of work in the market and at home. Despite the obvious importance of looking more closely at how people spend their non-work time, relatively little attention has been paid to describing its patterns and examining its determinants. A few studies have considered how the price of time affects the distribution of non-work time (Kooreman and Kapteyn, 1987; Biddle and Hamermesh, 1990); Aguiar and Hurst (2007) have charted secular changes in the distribution of non-market time in the United States, and Freeman and Schettkat (2005) have examined possible trade-offs between market and non-market work in a few countries. This line of inquiry has been limited by the relative paucity of available data sets. Until recently no country provided data on a continuing basis on how its citizens spend their time, and many have never provided such information. This absence of data has begun to change, and that change is what enables us to examine gender differences in the allocation of total work time.

The issue of gender differences in time use is important for a number of reasons. First, because the amount of work (and thus the utility from leisure) is one of the crucial arguing points in the “gender wars,” simply discovering new facts about it is important. Second, the determinants of those facts will allow us to infer how patterns of work by gender change as economies develop, thus shedding light on a crucial question in economic development. Third, by developing a new explanation for patterns of gender differences in the amount of total work, we may provide an impetus for using similar theories to examine other differences in the allocation of time. Finally, the facts we adduce and the theory we present to explain them can impose restrictions on a variety of economic models.

Our purposes here are to document in much greater detail a fact that has been essentially ignored by economists and that appears unknown to the public, and to offer and test some

explanations for that fact. In the next section we describe what we mean by market and household work and examine the gender breakdown of work at home and in the market using time-diary data from 27 countries. Whether the facts that we adduce in Section 2 are novel, and whether they are already widely known, are examined in Section 3. In Section 4 we consider some possible explanations of our findings and indicate which ones seem inconsistent with the results. This leads in Section 5 to the development of a theory of gender differences in total work time based on social norms. Section 6 examines some additional cross-country and micro evidence that is consistent with that theory and difficult to rationalize otherwise. The end result is a variety of facts and a theory that is not inconsistent with them, the totality of which might be used to inform how we model behavior in a variety of areas of economics, as we indicate in Section 7.

2 Gender Differences in Market and Home Work

In order to examine gender differences in work we need to devise general rules that allow activities to be classified as work. We follow standard practice and define market work as time spent for pay (or in unpaid household production for the market). We assume that people would not be working the marginal hour in the market if they were not paid, so that at the margin market work is not enjoyable (or at least is less enjoyable than any non-work activity at the margin). In the economics literature market work has generally been treated as the complement of the aggregate of all activities outside the market—implicitly all uses of non-market time have been assumed to be aggregable.¹

We count as household work those activities that satisfy the third-party rule (Reid, 1934) that substituting market goods and services for one's own time is possible. Such activities may be enjoyable (as may work in the market), even at the margin; but they still have the common characteristics that we could pay somebody to perform them for us and that we are not paid for performing them for ourselves. We define total work as the sum of time spent in market work

¹See Weiss (2009) for an excellent discussion of the development of philosophers' and economists' views on what constitutes work.

and household production. Note that we do not and cannot examine gender differences in the consumption value of the average or marginal minute of market or household production; all we do here is estimate, and then try to explain, differences in the total amount of time spent in productive activities.

We exclude all other activities from the definition of total work. Thus for our purposes it does not matter whether non-work time is used for sleeping, eating, religious activity, television-viewing, exercise, etc.: Our interest is in activities that might be defined as work and thus possibly be valued directly in the market.

Throughout this initial empirical section we define the aggregates of activities as similarly as possible across the countries under study. Respondents in these studies are given a time-diary for one or more recent days and asked to account comprehensively for all time during that day by time of day. The respondent either works from a set of codes indicating specific activities, or the survey team codes the descriptions into a pre-determined set of categories. Wide-scale time-diary surveys have been conducted for nearly 70 years (Sorokin and Berger, 1939).

No matter how extensive a set of activity codes is, each survey will have a different way of categorizing and aggregating what might seem like the same activity to an observer. Time diaries have the virtue of forcing respondents to report a time allocation that adds to 24 hours in a day. Also, unlike retrospective data about last year's or even last week's time spent working, while the time-diary information is necessarily based on recall, the recall period is only one day. The shorter recall period and the implicit time-budget constraint suggest that information on market work from time diaries is likely to be more reliable than the recall data on time use from standard household surveys; and, of course, time diaries provide information on non-market activities that is generally unavailable in labor-force surveys.

The aggregates from Australia, Germany, Italy, the Netherlands, Spain and the United States, and Israel that we use represent our own calculations from micro data in these sets of time

diaries. We appended our own aggregations based on published summaries of recent time-diary studies from other countries, including seven wealthy EU countries (Belgium, Denmark, France, Finland, Sweden, the United Kingdom, and Norway) and three transition countries (Estonia, Hungary and Slovenia) from Aliaga and Winqvist (2003); from various published summaries describing the results of time-diary studies conducted since 1992 in Canada, Ireland, Japan and New Zealand, and a set of sub-Saharan countries, Benin, Madagascar, Mauritius and South Africa (from Blackden and Wooden, 2006), and from Mexico and Turkey.²

The definitions of total work for each of the countries are shown in the Appendix. Obviously they are not identical across countries—but they are identical across gender within country. We cannot prove the absence of a systematic bias in the aggregations of results in each diary toward counting as work activities those performed especially by one gender or the other; but for our cross-country results to be biased would require systematic errors that are the same in most countries' methods of categorizing work activities.

Among the 27 countries that we study, the unweighted average of total work among women is 446.4 minutes per day (standard error = 8.6), while among men it is 421.7 minutes (s.e. = 8.9). Although these averages differ significantly from one another, if we restrict the sample to the 14 wealthy (2002 real GDP/capita above \$15,000, from Heston *et al*, 2002 and 2006) non-Catholic countries the averages are 440.1 (s.e. = 7.4) and 431.4 (s.e. = 7.5) respectively. The gender difference in total work is quite small and not significantly different from zero in the richer countries. We will refer to this striking outcome henceforth as the **iso-work** phenomenon. Given the fixed length of the day, this finding implies that there also is gender equality in the total of non-work time consumed.

The scatter diagram in Figure 1 compares men's and women's total work in the 27 countries. The steepest line shows what women's total work would be if it were identical to

²While a few of the underlying micro data sets are available, most are not, thus necessitating our reliance on published sub-aggregates in constructing the aggregate of total work.

men's. We then estimated a regression relating the amount of total work among women to that among men. To accommodate the observation that women's total work always exceeds men's in Catholic countries, we included an indicator for this religious background (equaling one in six countries). The regression results (coefficient estimates and standard errors) are:

$$\text{FemaleTotalWork} = 134.05 + 0.73\text{MaleTotalWork} + 17.72\text{Catholic}, N=27, R\text{Bar}^2 = 0.590 .$$

$$(51.54) \quad (0.12) \quad (13.36)$$

(The regression line through the non-Catholic points is the lower of the two parallel lines in Figure 1; the line fitting the points describing Catholic countries is the upper parallel line.) We can reject the hypotheses that the intercept is 0 and that the slope on MaleTotalWork is 1, as well as the joint hypothesis.³ Nonetheless, the slope of the relationship between total work by gender is not that much different from 1; and, as we showed, the averages for rich non-Catholic countries do not differ significantly statistically or economically.

To examine the role of economic development, Figure 2 shows a scatter of the difference in average minutes per day of female over male total work time and real GDP/capita, along with a line fitting these points. The scatter and fit suggest either that economic development is highly positively correlated with gender equality of total work, or that today's rich non-Catholic countries have always had a different culture along this dimension from today's poor countries and from Catholic countries. Furthermore, the relationship appears to be nonlinear, with gender equality being approached or even reached at a sufficiently high level of personal income.

Clearly, based on these results we cannot claim that this remarkable gender equality in total work holds at all times and places. Our results also show that it does not hold in middle- or lower-income countries; and Haddad *et al* (1995) suggest similar findings for other developing African economies, as do Goldschmidt-Clermont and Pagnossin-Aligisaks (1995) for Bulgaria in 1988. The evidence, however, does suggest strongly that iso-work characterizes average

³The statistic testing the joint hypothesis is $F(2,24) = 8.18, p=.002$.

household behavior and labor markets in rich countries generally and is positively associated with national differences in real incomes.

3 Novelty and Knowledge

The iso-work fact has been demonstrated by several sociologists. Robinson and Godbey (1999) employ data from a UN report (Goldschmidt-Clermont and Pagnossin-Aligisakis, 1995) to show that this fact describes the average of (recall and time-diary) data from 14 countries from the 1980s and early 1990s; and Gershuny (2000) shows that it roughly characterizes the two averages over a similar sample of data sets covering the 1960s through the mid-1990s. No study has demonstrated it using data sets that were as well harmonized as those that we used here, nor has one shown how closely it describes average outcomes in individual countries.

The fact is thus not new in the sociology literature, although it appears not to have been studied in the economics literature. The difficulty, however, is that it has been swamped by claims in widely circulated ethnographic studies (Hochschild, 1997, and earlier work) based on a few households that women's total work significantly exceeds men's. Indeed, even sociologists who have demonstrated it (e.g., Mattingly and Bianchi, 2003, for the United States, and Bittman and Wajcman, 2000, for several countries), downplayed it to focus on showing that women's work is more onerous than men's, and why women's leisure provides less pleasure.

With all this evidence demonstrating gender iso-work, one wonders whether the fact that we have demonstrated is well known among economists, other social scientists and the general public. To examine this issue, we designed a survey that asked the following question:

“We know that American men (ages 20-75) on average work more in the market than do American women. But what is the difference between men's TOTAL WORK (in the market and on anything that you might view as work at home) and that of women? Without consulting any books, articles or raw data, PLEASE PUT AN X NEXT TO THE LINE BELOW THAT YOU BELIEVE TO BE THE CLOSEST APPROXIMATION TO THE CURRENT SITUATION IN THE US.”

Respondents were allowed nine possible responses, ranging from a 25 percent excess of female total work, to symmetry around equality, to a 25 percent excess of male total work.

In August 2006 we emailed this survey to three groups: 1) 663 labor economists affiliated with a worldwide network of such researchers. The web-based survey allowed us to distinguish respondents who had spent at least six months in the U.S. from those who had not; 2) 255 elite macro and public finance economists, members of a mostly American network of such researchers; and 3) 210 faculty members and graduate students in a leading American sociology department. The first and third groups received follow-up emails three weeks after the initial survey. Also, early in September 2006 we asked the same question of 533 students in an introductory microeconomics class. Using the information on location in the first group, we thus have five separate sets of responses. The response rates varied, but there is no reason to believe that non-respondents were less well-informed about the facts than respondents.

The results of these surveys are shown in Table 1. The majority of respondents in each of the five groups believe that American women perform at least five percent more total work than men. Assigning half the respondents who state that there is equality to this category, we convincingly reject the null hypothesis that the proportions stating that men work less or women work less are equal. Indeed, even if we assign all those stating that there is equality to the “men work more” group, this null hypothesis is rejected in some of the samples.

Samples of sociologists, experts in labor economics, leading economists and the public believe that women work more in total than do men. Indeed, the results from the freshmen (half of whom are women) look very similar to those from labor economists and, in turn, to those from macroeconomists, even though all three populations have substantially different fractions of women. Perhaps the only consolation is that the distance between opinion and fact is less within the groups of economists than it is among the sociologists. Despite our demonstration of gender equality of total work in most rich countries using current time-diary data, and despite demonstrations using time-diary and recall data of this general fact by several sociologists, all

groups considered here appear ignorant of the reality. It may be that most people are aware of the underlying differences and that our samples are unusual; and, as is always true with eliciting information about beliefs, perhaps the respondents to our surveys were thinking about something other than the sheer quantity of work—perhaps how onerous it is. Nonetheless, this evidence suggests as strongly as survey evidence on subjective beliefs can that people believe women's total work exceeds men's in rich countries.

4 What Fails to Explain the Facts?

Assuming substitution effects dominate income effects, economic theory predicts that a rise in women's relative wage (i.e. the gender wage gap) will lead to more work in the market by women relative to men. The impact of this increase on the relative amount of home work should be in the opposite direction, so that the effect of a change in the gender gap on the relative amounts of total work is ambiguous. Unless, however, additional market work is offset one-for-one by reduced home work, a rise in the female relative wage should raise women's relative total work.

To examine this possibility, we use Polachek and Xiang's (2006) estimates of the gender wage gap. In particular, for 19 of the 27 countries on which we have recent time-diary data they produced estimates of the difference between the logarithms of the medians of the distributions of males' and females' wages. Using these data, in the first two columns in Table 2 we present least-squares estimates of equations describing female-male differences across countries in market and total work as a function of the gender pay gap. The results on market work are consistent with an upward-sloping relative supply curve of labor to the market. The market work effect, however, dominates the household work effect, so that we find that the female-male gap in total work is also positively related to the female-male wage ratio.

These findings are not affected by the inclusion of real GDP per capita, as the estimates in Columns (3) and (4) show, nor are they affected by adding the indicator for Catholic countries. That the GDP variable is only marginally statistically significant, whereas Figure 2 suggested a strong negative relationship with a diminishing slope, arises from the exclusion of many of the

poorer countries where wage data are unavailable. Higher relative wages among women lead them to work relatively more in the market, less at home, and more in total. Despite the quality of the estimates, the equation in Column (6) describes below half of the variance in the gender difference in total work across countries. The difficulty is that in 14 of these 19 countries gender differences in total work are clustered within five percent of equality, while the gender wage gaps in these data range from 0.07 to 0.69. Something, not equality in relative wages or differences in per-capita incomes, is causing the pervasive absence of gender differences in total work.

Taking a different perspective on these results, one might follow the literature on household behavior (e.g., Lundberg and Pollak, 1996) and view the gender wage gap as indicating differences in power in the household, as it would be regardless of whether one views spouses' behavior as described by a unitary or a collective model. By this criterion, we should expect that where female-male relative pay is lower we will observe men working relatively less in total.⁴ The estimates in Table 2 imply the opposite result. Where one might infer that men have more bargaining power, as measured by relative wages, their total work is in fact greater.

A second possible explanation for some of these facts is that husbands and wives pay attention to each other's labor and leisure, so that we observe gender equality at the means in rich countries because most adult men and women are married. To examine this possibility in the aggregate, in Table 3 we present means of market work and home work by gender and marital status for the United States in 2003 and Germany in 2001/02, two of the few countries for which we have micro-data. While the female-male gap in total work is higher among unmarried adults,

⁴Consider a large class of models in which spouses derive utility from consumption, which is public and joint to the couple, and separably from leisure, which is the time-budget complement of total work time. The couple must produce a fixed amount of public consumption at home using a constant returns production function of each household member's time, without prejudice to the relative efficiency of the man or woman. Assume that the joint family decision is to maximize a weighted sum of utilities of the two household members. The solution to this two-stage problem is a labor supply rule which implies iso-work if and only if each utility weight equals the percentage deviation of the respective gender wage from its average. Thus the greater the excess of male over female wages, the lower men's total work relative to women's.

in the United States it varies across marital status within 5 percent of equality. In Germany the gap is larger among unmarried adults, but still not huge.

An explicit test of the notion that gender iso-work is generated by husbands and wives focusing on each other's work effort as part of marriage can be conducted by examining inter-household dispersion in the within-household gender total work gap. Because only one person per household was sampled, this examination is not possible for the U.S. in 2003, so instead we use the much smaller 1985 U.S. Time Use Survey, which collected data on both spouses on a single day. For the 2001/02 German data this is easy, as diaries were collected from both spouses on three separate days, which we average for each spouse. As another comparison we examine data from the 1992 Australian survey, which we summarized in Figure 1, averaged over the two days on which each spouse kept a diary.

Figures 3a-3c show the frequency distributions of the differences within households between the average daily total work of wives and husbands in the U.S., Australia and Germany. While the distributions are symmetric around means of zero, the implied dispersion is large in each case. Indeed, regressions within each country of the wife's total work time on the husband's explain only 9 percent of the variation in the U.S., 29 percent of the variation in Australia and 35 percent in Germany.⁵ While wives work more in total when their husbands work more, within-couple covariation describes only a small part of the variance in spouses' total work time. We do find evidence of complementarity of spouses' total work (and thus of leisure), but most of the dispersion in intra-household differences in total work remains.⁶ This evidence is inconsistent with the assertion that the iso-work phenomenon stems from the alignment of behavior within a couple.

⁵If we examine differences in daily total work, and in Germany two-day averages, the unexplained variances are higher than those discussed in the text; but the gain to aggregating over more days drops rapidly as the number of days aggregated is increased.

⁶See Hamermesh (2002) for evidence of the importance of both the quantitative and timing aspects of complementarity of spouses' market work.

5 Social Norms in the Theory of Total Work

What kind of mechanism could coordinate total time spent on market work and secondary activities across males and females on average, regardless of whether they are married or unmarried? The simplest coordination device that equalizes total work across agents is a *social norm for leisure* that serves as focal point for the determination of total work. Peer pressure or a strong desire to conform to a common social norm for time allocation mutes market incentives and weakens the impact of individual tastes. As a result, time use becomes more similar across individuals.⁷ If the social norm is strong enough to drive the agent to conform fully, we obtain the iso-work result we observe in the data.⁸ Alternative explanations of the iso-work fact are, of course, possible; but all must involve, in one way or another, an interplay between social factors and individual tastes. This interplay might lead, of course, to multiple equilibria. We relegate multiplicity issues to the Appendix in order to focus on iso-work.

5.1 Baseline model

Imagine that, *in the absence of a social norm*, consumers maximize the linear-quadratic utility function

$$C - (1/2\epsilon)(1 - L)^2 \tag{1}$$

subject to constraints

$$C = \Omega + wH, \tag{2}$$

$$H + L = 1, \tag{3}$$

where C and L denote consumption and leisure, w is the wage rate, Ω represents non-labor income, the parameter $\epsilon > 0$ is an (inverse) index of the marginal disutility of work, and without loss of

⁷For a survey of social norms and economic theory, see Elster (1989). Social norms have been studied, among others, by Akerlof (1980), Jones (1984), Cole et al. (1992), Kandori (1992), Young (1996), Lindbeck (1997), and Lindbeck et al. (1999).

⁸In this simple story, total conformity only occurs if the desire to conform is infinitely strong. The literature on conformity (Bernheim, 1994) has sought ways to obtain identical behavior without assuming an infinite cost of deviation.

generality the amount of available time is normalized to 1.⁹ Optimal leisure is then

$$L = 1 - \epsilon w. \quad (4)$$

We call this the agent's *intrinsic* leisure optimum. It is determined by private incentives, prices and budget constraints.

5.2 Linear-Quadratic Leisure Norm

Now suppose that there is a social norm that influences, but does not mandate, individual leisure. We mean by this that agents can choose the extent to which they adhere to the norm, and that they balance the marginal costs and benefits of deviating from it. The cost of deviating may stem from guilt (an internal psychological process) or shame (an external peer pressure or a reputational mechanism). The benefit of deviating results from the joy of following one's own unbridled inclinations that in general differ from the norm.

Formally, assume that there is a quadratic cost of deviating from the leisure norm L^* , and parameterize the strength of the social norm by the coefficient $\phi \geq 0$,¹⁰ so that the utility function becomes

$$C - (1/2\epsilon)(1 - L)^2 - (\phi/2)(L - L^*)^2. \quad (5)$$

Optimal leisure is

$$L = \alpha(1 - \epsilon w) + (1 - \alpha)L^* \equiv L(w), \quad (6)$$

with the weight α , between 0 and 1, given by

$$\alpha = \frac{1}{1 + \phi\epsilon}. \quad (7)$$

Intuitively, the social norm pulls optimal leisure away from the intrinsic optimum $1 - \epsilon w$ and towards L^* . The coefficient α will be small, and optimal leisure will be close to the norm, if the social norm is strong (ϕ large) or leisure is not too wage inelastic (ϵ large). Higher wages, holding α

⁹We impose the restriction $\epsilon > 0$ to exclude backward-bending labor-supply curves. Here, and in what follows, we assume that non-negativity constraints on consumption and leisure are satisfied. In particular, we assume that the wage is always below $1/\epsilon$ to avoid corner solutions at $L = 0$.

¹⁰The strength of the norm for an individual may depend on the number of people who have adopted it. We examine this possibility below.

constant, increase the distance between L and L^* by making it more costly to deviate from the intrinsic optimum.

Now suppose male (m) and female (f) wages differ, but that there are no within-gender wage differences, and therefore no within-gender leisure difference. Assume also, to simplify, that the wage sensitivity of leisure (α) is the same for both sexes.¹¹ As a result, the average (and individual) leisure of agents of gender i is simply $\bar{L}^i = L(w^i)$, and aggregate *leisure gap* is

$$\bar{L}^m - \bar{L}^f = L(w^m) - L(w^f) \tag{8}$$

$$= -\alpha\epsilon(w^m - w^f). \tag{9}$$

Since α is decreasing in ϕ , equation (8) implies iso-work—a leisure gap close to zero—if the norm is very strong. In the limit, $\lim_{\phi \rightarrow \infty} (L^m - L^f) = 0$. In words, a very strong norm mutes the effect of differing male and female wages on leisure, thereby leading to iso-work by equalizing male and female leisure in the aggregate.

This result highlights an essential ingredient of any norm-based explanation of the iso-work fact: the fact that the fraction of men and women that share a given norm is identical. We call this feature the *gender-neutrality* of norms, and we will show below that it is crucial for iso-work to emerge in the presence of many norms and within-gender wage heterogeneity. In the current example, it is because all males and all females have the same leisure norm that a larger ϕ eliminates the differences between male and female leisure. Were the fraction of men and female who adopt the norm L^* different, we would, *ceteris paribus*, observe different male and female average leisure even when $\phi = +\infty$. Hence the fact that total work is essentially invariant to gender in high-income countries (but less so in poorer economies) suggests, if the social norm story is correct, that a fundamental change of norms takes place in the process of economic development: gender-neutral, or gender-blind norms replace gender-specific social reference.¹²

¹¹This last assumption, which is of course at odds with estimates of labor supply elasticities for males and females, can easily be relaxed.

¹²Note that no causal statement is being made here. One can easily write models in which gender-specific norms cause economic backwardness, and models in which competition and development cause gender equality.

5.3 Accounting for Within-Gender Heterogeneity

Although it provides us with an important insight, the small model we have just outlined is not sufficient to rationalize all the facts. The empirical difficulty we face is that the iso-work fact coexists with significant within-gender (and more generally within-group) heterogeneity of leisure. This is inconsistent with the simple story told above, because as $\phi \rightarrow +\infty$ the labor supply of each individual, whether male or female, converges to the common, gender-neutral norm L^* regardless of the wage.¹³ As a result, while a strong norm bridges the gap between male and female leisure, it also eliminates any within-gender heterogeneity of leisure.¹⁴ This unpleasant feature can be avoided by allowing for non-gender based *social clusters*, or multiple social norms.

5.3.1 Social Clusters

Imagine that each gender can be stratified into social clusters that are defined by the *relative* position in the wage distribution. Assume, for instance, that males and females above their gender's *median* wage share a common leisure norm, and that there exists another leisure norm for males and females below the median wage. Note that we could just as well cluster agents according to the color of their eyes, the month in which they are born, or the neighborhood in which they live. The crucial assumption, as hinted above, is that the clusters are defined by gender-neutral characteristics: the fractions of men and women above the median wage of their respective gender are identical, and so are (presumably) the proportions of men and women who have blue eyes, are born in December, or live in Austin, TX.¹⁵

Suppose, to be more formal, that there are *two* leisure norms L_h^* and L_l^* ,¹⁶ and that an individual, male or female, adopts norm L_h^* (resp., L_l^*) if he/she is in the upper q -th (resp. lower upper $1 - q$ -th) percentile of the respective gender's wage distribution. Define the threshold wages w^i by

$$1 - F^i(w^i) = q, \tag{10}$$

¹³This is also true if ϵ , the sensitivity of leisure to the wage, differs across sexes.

¹⁴Note that the theory of conformity developed by Bernheim (1994) to explain why people with different intrinsic preferences behave identically suffers, from the perspective of iso-work, from the same weakness as the linear-quadratic model with one norm $\phi = \infty$: it wipes out all within-group heterogeneity.

¹⁵By contrast, social leisure norms defined in terms of position of the wage above or below some arbitrary *levels* (i.e., a leisure norm for “high” wage males and females, another one for “low” wage ones) will be in general gender-biased, as the proportions of males and females adopting a given norm will differ unless the separating levels happen to coincide with median wages.

¹⁶The analysis can be generalized readily to many norms.

where $F^i(\cdot)$, $i = m, f$, is the cumulative distribution of wages for gender i . Thus, men with wages above (below) \mathbf{w}^m and women with wages above (below) \mathbf{w}^f adopt norm L_h^* (L_l^*). Assume that the strength ϕ of the social norm is the same for all individuals. Leisure of an agent of wage type j with wage w is simply, as before,

$$L_j(w) = \alpha(1 - \epsilon w) + (1 - \alpha)L_j^*, \quad (11)$$

so that the average leisure of agents of gender i and wage w is

$$\bar{L}^i = \int_{w < \mathbf{w}^i} L_l(w) dF^i(w) + \int_{w > \mathbf{w}^i} L_h(w) dF^i(w) \quad (12)$$

$$= \alpha(1 - \epsilon \bar{w}^i) + (1 - \alpha)[(1 - q)L_l^* + qL_h^*]. \quad (13)$$

We immediately conclude that the average leisure gap between men and women is

$$\bar{L}^m - \bar{L}^f = -\alpha\epsilon(\bar{w}^m - \bar{w}^f). \quad (14)$$

This is the same formula we obtained with a single social norm. As before, the leisure gap goes to zero and the iso-work fact holds asymptotically when the social norm becomes infinitely compelling ($\phi \rightarrow \infty$, so that $\alpha \rightarrow 0$). However, the existence of many social clusters (delineated by categories that are orthogonal to gender) ensures that *within-gender* heterogeneity of leisure does not shrink to zero as ϕ becomes large. Indeed, this orthogonality condition is necessary for the existence of iso-work within any particular sub-group of the population.

5.3.2 Even More Heterogeneity

In Section 4 we demonstrated significant heterogeneity of total work, even within couples. One way to capture this aspect of the data is to define higher dimensions of clustering based on other characteristics of agents, and to repeat the foregoing reasoning for this finer partitioning of the population. By doing so—provided of course the resulting categories are uncorrelated with sex—we could again replicate the iso-work fact yet generate as much within-gender heterogeneity as desired by making each social norm increasingly compelling. Of course, we would still find that within-category heterogeneity would go to zero, but this would not be much of a problem anymore as the categories

would be arbitrarily fine.

Alternatively, within-category heterogeneity as norms become more binding may reflect idiosyncratic heterogeneity in the population. This heterogeneity could stem from different tastes, or from a noisy individual observation of the societal leisure norm.¹⁷ To illustrate how this line of reasoning would play out in our setup, return to the first of our models with one norm L^* for all, identical wages for all members of a given sex, and a different wage for male and female workers. Imagine that individual k observes the norm with some measurement error λ_k , believing that the desirable norm is $L^* + \lambda_k$ instead of L^* .¹⁸ As a result, optimal leisure for that individual becomes

$$L_k = \alpha(1 - \epsilon w) + (1 - \alpha)(L^* + \lambda_k), \quad (15)$$

with α defined exactly as above. Hence $L_k \rightarrow L^* + \lambda_k$ as $\phi \rightarrow \infty$ (and $\alpha \rightarrow 0$) regardless of the wage, i.e. regardless of whether one is male or female. Now suppose further that measurement errors are idiosyncratic in the sense that the λ 's average to zero for each sex.¹⁹ Then it is straightforward to show the leisure gap is zero, and the iso-work fact holds exactly when $\phi \rightarrow \infty$ —in spite of the fact that each agent ends up taking a different amount of leisure due to an idiosyncratic perception of the norm.

5.4 Accounting for Variations in Total Work

The data presented in Section 2 make it clear that, although total work is strikingly equal across men and women, it does vary, sometimes substantially, across countries, region and over time. Since we have attempted in the previous section to rationalize the iso-work fact by social norms by arguing that they serve as a coordination device between male and female total work, we must also explain how norms can vary. This is most simply done by endogenizing the leisure social norm.

Let us return yet again to our simplest model of social norms: men and women have the same preferences, they face a gender-specific wage, there are no within-gender wage differences, and men and women adopt a common leisure norm L^* . Remember that in that model average male and female

¹⁷As we do not wish to transform the quest for a theoretical explanation of the iso-work fact into a futile data-fitting exercise, we prefer the second interpretation, which is potentially falsifiable, to the first, which increases the number of unobservable parameters.

¹⁸For example, an individual of type k has utility function $C - (1/2\epsilon)(1 - L)^2 - (\phi/2)[L - (L^* + \lambda_k)]^2$.

¹⁹This leaves open the possibility that females and males perceive the social norm with different precision.

leisure are given by

$$\bar{L}^m = \alpha(1 - \epsilon w^m) + (1 - \alpha)L^*, \quad (16)$$

$$\bar{L}^f = \alpha(1 - \epsilon w^f) + (1 - \alpha)L^*. \quad (17)$$

Now close the model by assuming that the gender-neutral norm L^* reflects *average leisure* across males and females in society. Since there are equal proportions of men and women, in equilibrium we have

$$L^* = \frac{1}{2}(\bar{L}^m + \bar{L}^f) = L^*. \quad (18)$$

Combining the last three equations and solving for L^* , we conclude that the *equilibrium social norm for leisure* is simply

$$L^* = 1 - \epsilon \bar{w}, \quad (19)$$

where

$$\bar{w} = \frac{w^m + w^f}{2} \quad (20)$$

is the average wage in the overall (male *and* female) population. The equilibrium social norm is independent of the strength of the norm, but it is negatively affected by the average wage rate \bar{w} , with a response coefficient that depends on the sensitivity ϵ of individual leisure to the wage. Whenever these magnitudes change, across countries or over time, the social norm for leisure varies. There is no reason to expect it to remain to be constant.

5.5 Accounting for the Relationship between GDP Per Capita and the Female-Male Total Work Difference

We have argued above that female-male differences in total work are negatively related to GDP per capita. A social norm theory of leisure can deliver this fact in at least three, not mutually exclusive three ways. The first relies on the link between economic development and the increased gender-neutrality of social reference groups. The second, which is slightly more *ad hoc*, assumes that the cost of deviating from a social norm is positively related to the wage. A third approach, inspired by Goldin (1995), postulates that women are socially stigmatized when they participate in some types of market activities.

First, a model of social clusters can account for the reduction in the female-male total work difference as GDP per capita grows provided economic growth is positively correlated with the adoption of gender-neutral reference groups. Suppose for instance that at low income levels there are two leisure reference groups: one for men, and one for women, each with a different (gender-specific) leisure norm. This might be due to a taste for discrimination, for example, which is correlated with income level. Then, trivially, iso-work does not hold at low income levels. If gender-defined social clusters are replaced by gender-neutral reference groups as income rises (e.g., at quantiles of income distributions), then development will be associated with a convergence of the total work difference across genders to zero.

An alternative, possibly complementary explanation relaxes the assumption that deviating from the norm entails a utility cost that is independent of the level of the individual's wage. Let us consider in the simple one-norm model what happens if people get harassed when they deviate from the norm. That is, imagine that, instead of suffering a direct utility loss as envisaged above, deviants lose time fending off their critics, mending their reputation, or battling inner guilt feelings at the cost of time available for work or leisure. Namely, they solve:

$$C - (1/2\epsilon)(1 - L)^2 \tag{21}$$

subject to constraints

$$C = \Omega + wH, \tag{22}$$

$$L + H + \frac{\phi}{2}(L - L^*)^2 = 1. \tag{23}$$

It is straightforward to show that the solution to this problem is formally equivalent to that of the utility-loss model, with the parameter ϕ replaced by ϕw . In other words, the “harassment” model is just the utility loss model with a cost of deviation proportional to the wage. Therefore, adapting equation (6), we conclude that optimal leisure in this model is

$$L = \alpha(w)(1 - \epsilon w) + [1 - \alpha(w)]L^*, \tag{24}$$

with the weight $\alpha(w)$ given by

$$\alpha = 1/[1+\varphi\epsilon w]. \quad (25)$$

At a low wage or level of development (w close to zero), the weight $\alpha(w)$ is close to 1, so that the intrinsic optimum $1 - w$ is the main determinant of leisure. At a high wage or development level (w high), and given the parameter φ , the weight $\alpha(w)$ approaches zero, and the social norm becomes the sole determinant of optimal leisure. As the value of time increases, so does the cost of deviating from the norm, resulting in a smaller deviation from the norm.

A final possibility is that, as in Goldin (1995), there is a social stigma attached to female participation in some market activities. Goldin (1995) assumes that blue-collar, but not white-collar work by a woman engenders a fixed utility loss S . For our purpose, imagine a simpler scenario in which **any** female market activity is stigmatized, and in which there is no social norm beyond this stigma. If a woman works in the market at wage w , her utility is, according to the baseline model of Sub-section 5.2, $(\Omega + \epsilon w^2/2) - S$, or the difference between the utility stemming from the intrinsic optimum described in (4) and the utility loss entailed by market participation that arises from the stigma. By contrast, staying home to conform to the way women are stereotyped provides utility Ω . Staying home is optimal as long as $\Omega > (\Omega + \epsilon w^2/2) - S$, or $\epsilon < [2S]/w^2$. Development, and the concomitant rise in wages, thus reduces the impact of gender stereotypes on behavior. In that respect, development makes men and women behave, *ceteris paribus*, in increasingly similar ways.²⁰

6 Some Evidence on the Role of Social Norms in Iso-Work

The theory developed in the last section is not easy to test directly. It does, however, provide some guidelines for analyzing the data. First, any test based on arbitrary sub-groups in which the

²⁰An alternative justification for the correlation between the strength of social norms and GDP can be inferred from experimental work conducted in rural Africa (Ensminger, 2004). One reason for the strength of this correlation may be that the physical demands of market work during the early stages of industrialization penalize women so that they must work more hours in order to supply the same amount of productive effort as men.

gender ratio differs (e.g., groupings by age, absolute earnings, etc.) is excluded by the theory. Second, conditioning on outcomes (e.g., examining differences by gender among paid workers) both violates the orthogonality condition and confounds effect with cause. We can, however, perform a number of indirect tests of the consistency of social norms with observed behavior. None of these can conclusively validate the theory to the exclusion of others; but each offers a chance to examine whether the theory can be shown to be inconsistent with a variety of facts gleaned from various data sets.

If the notion of social cluster norms is correct, differences in total work across various candidate references groups will be large compared to gender differences within clusters. Consider first cutting the data by educational category. In the 2003 U.S. data we divide the adult population into those with fewer than 12 years of school, 12 years of school, some college, and college or more. In the German data we create four categories: *Volksschule/Hauptschule* (basic), *Mittlere Reife/Realschule* (high school), *Fachoberschule/Fachabitur* (vocationally qualified), and *Abitur* (university preparatory).

Table 4 shows average minutes of total work by gender for each of the education categories in the U.S. and Germany. In both countries, gender differences in total work within education categories are small, with the highest being the 5 percent excess of female over male total work among the least educated Americans. Differences across categories in total work independent of gender are significantly larger: In the United States the gap between the highest and lowest education categories in the average amount of total work is 39 percent, while in Germany it is 13 percent. Clearly, gender differences are tiny compared to those resulting from differences in educational attainment.

Similarly, the data could be cut by region. To the extent that there are inter-regional cultural differences, we might expect different norms regarding total work across regions, even though gender differences within regions are small. The possibilities for examining this notion are limited by sample size in both data sets. Also, confidentiality restrictions on the German data

prevent us from obtaining a finer geographic breakdown than West and East. With these data limitations, we divide the U.S. sample into the South and non-South Census regions, and the German data into West and East.

Averages of total work by gender within geographic area are shown in Table 5. Notice first that within-region differences in total work by gender are not large. While those in the South and within each German region are statistically significant, none exceeds 3 percent. Among Southern women, total work is over 5 percent below that in the rest of the nation, while among Southern men it is 3 percent below. For Germany we observe a qualitatively similar outcome: West-East differences in total work are 4 percent among women and 3 percent among men. The contrast between inter-regional differences in average total work and within-region differences by gender is consistent with the notion of clustering on norms, although the differences are not as stark as those observed when we cut the data by educational attainment.

An additional piece of evidence based on aggregated data asks, without any claims of causation, whether attitudes about gender roles are related to gender differences in total work. To examine this relationship we use data collected in various years by the World Values Surveys (WVS). Respondents in many countries were asked whether they agreed with the statement that men should have more right to scarce jobs than women. Taking averages of these data for each country for the most recent year before the time-diary survey, we present in Figure 4a a scatter diagram relating them to the female-male difference in total work for 21 of the 27 countries used in Figure 1 for which they are available. The scatter and the highly significant relationship between the gender total work difference and this attitudinal variable that we might interpret as representing beliefs in male pre-eminence suggest that, where the expressed norm about the labor market favors men, women perform a greater share of the total amount of work.

One might argue that this diagram merely reflects generalized cultural differences rather than specific attitudes about gender work roles. To examine this possibility we use a general measure of attitudes toward work in the WVS, the fraction of respondents agreeing that it would

be unfortunate if there were less emphasis on work in the future. The scatter of this variable and the gender difference in total work is shown in Figure 4b, along with the regression relating the two. The fit is much worse than with the variable measuring attitudes about gender work roles. Taking this line of argument one step further, we obtained an attitudinal measure from the WVS that reflects culture but is unrelated to attitudes about work—the fraction of respondents stating that they are very proud of their nationality.²¹ The scatter and regression of this variable and the gender difference in total work presented in Figure 4c show no relation between the two.²² While we do not claim causation from these measures of attitudes to behavioral outcomes, the exercise does suggest a link between those outcomes and specific attitudes about gender roles in work.

Evidence on norms and their evolution in this context can be gleaned from the behavior of immigrants relative to that of their descendants (see Fernandez, 2007). To examine this possibility we first calculated total work time using diaries from the American Time Use Survey 2003-2006, which contain over 6000 immigrants. We were able to identify immigrants from 26 countries for which the World Values Survey included responses to the three questions used in the analyses of Figures 4a-4c. For immigrants from these 26 countries we then related the average total work of all married women in the immigrant group to the average total work of all men in the group and, crucially, sequentially to each of the three variables used in Figures 4a-4c. To account for possible differences in life-cycle patterns in labor-force participation we also controlled for the average ages of husbands and wives in each immigrant group. The regressions are weighted to account for the number of immigrant wives from each origin country who are included in the ATUS.

²¹All the data can be downloaded from <http://www.worldvaluessurvey.com/>. The questions are: 1) “Do you agree or disagree with the following statement, When jobs are scarce, men should have more right to a job than women.” 2) “Please tell me, if it were to happen, whether you think it would be a good thing, a bad thing, or don’t you mind: Less importance placed on work in our lives.” 3) “How proud are you to be [Nationality]?”

²²If we replace national averages of attitudes in each of these scatters with gender-specific national averages the results are hardly unchanged, as the correlations of averages of female and male attitudes in each case exceed 0.9.

The results of this estimation are shown in the first three columns of Table 6. The crucial finding is in Column (1)—among immigrants from those countries where natives believe more strongly that men should have priority for scarce jobs women perform more total work than do men. Both across countries and among immigrants to the U.S. the iso-work phenomenon is affected by attitudes toward gender roles in work. The estimates presented in Columns (2) and (3) demonstrate that this result is specifically due to attitudes about the relation between gender and work: When we examine the impacts of attitudes about work generally, or the effect of an unrelated cultural attitude, we find no impact on wives' total work compared to husbands'.

The ATUS data allow us to perform similar calculations for second-generation Americans. We can identify such individuals from 25 origin countries, and the final three columns of Table 6 list estimates for them constructed exactly like those for immigrants. There is absolutely no relationship between differences in any of these three attitudinal measures in their parents' home countries and the female-male difference in total work. The social norms about gender work roles that exist in other countries clearly disappear, at least at the national-origin level, after one generation in the U.S.

We can go further toward examining the unique causative role of norms in generating gender differences in total work using the micro data on couples in Australia and Germany that underlie the histograms in Figures 3b and 3c. For each couple (3080 in Germany, 1966 in Australia), we initially regress wife's total work (averaged over three diary days in Germany, two in Australia) against husband's, with the results shown in Columns (1) and (5) of Table 7. Following a now-substantial literature on the role of peer effects in behavior (see Borjas, 1992, for an early example), in Columns (2) and (6) we add to these regressions variables measuring the average work of wives in the particular wife's education group (four in Germany; three in Australia), age group (<40, 40-54, 55+) and region (West and East in Germany; New South Wales, Victoria, Queensland and other jurisdictions in Australia). Columns (3) and (7) present the same regressions with controls for the wife's own demographic and family characteristics. In

both countries the addition of these peer outcomes improves the ability of the equations to describe the wife's total work conditional on her husband's. Except for age for Germany in Column (3), peer outcomes of similarly-situated wives have significant impacts on the total work of individual wives.

One might argue that these results merely indicate the importance of Manski's (1995) reflection problem. We cannot demonstrate causation conclusively, but indirect evidence suggests that our results do not arise solely from the reflection of one's own behavior. First, while we cannot construct finer education and region groups in these samples, we can disaggregate wives' peer outcomes more finely by age—into seven age groups in the larger German sample and six in the Australian sample. Re-estimating the equations in Columns (2), (3), (6) and (7) using this finer distinction on the age peer effects, the adjusted R^2 increases in all cases, by about 0.004 for the estimates for Germany and by 0.001 in the estimates for Australia.

This first test does not avoid the reflection problem and indeed might arguably exacerbate it. The problem is at least directly avoided if we randomly partition the German and Australian samples into halves, calculate peer averages for one half-sample and include them in regressions like those in Columns (3) and (6) based only on the other half-sample. The results of re-estimating the equations on the second half-samples differ little from those shown in the Table. These similar results cannot be based on reflections, as the half-samples are different. Individuals in the half-samples may, however, be responding to their own unmeasured common characteristics rather than to the behavior of their peers. There is no way of circumventing this potential difficulty completely, just as in the larger peer-effects literature extricating common effects from responses to peers' outcomes is exceedingly difficult. In the case of Germany, however, we can probe a bit further by including in the regressions describing wives in the 2001/02 sample peer outcomes based on wives on whom time diaries were obtained in a nearly identical survey in 1991/92. Those results are presented in Column (4) of Table 7 and should be compared to the results in Column (2). While the impact of wives' peers' education has been

vitiated when we include lagged values, the impacts of peer outcomes in the same age group and region, although attenuated, are still highly significant statistically.

7 Conclusion: The Importance of the Iso-Work Phenomenon for Economic Modeling

In Section 5 we showed that the iso-work fact and its link to economic development place tight constraints on the modeling of labor supply behavior by gender. Any nontrivial gender-neutral model of labor supply must rely on the existence of strong cluster norms which coordinate behavior, or rely on implausible mean-preserving transformations of underlying distributions which are in turn unlikely to be common across gender. Consequently, iso-work gives rise to a number of conundrums for economic models that rely on work-leisure choices to characterize economic behavior in both the short- and long-run.

The first implication is linked to long-run economic development. Our evidence documents a convergence of total work across gender with GDP per capita. We show in Section 5 that this convergence can derive either from increasingly gender-blind assignment to reference clusters with strong norms, or from a convergence of gender wage-offer distributions to a common one. The past half century has also seen secular, albeit slow convergence in gender wage differentials. These two phenomena are probably related, but what is their source? Has technical change augmented female market production relative to that of men? Is technical change in home production generally labor-saving (see Greenwood *et al*, 2005)? More likely, how have interactions of these two types of innovation combined to generate the convergence in total work and the returns to market work? Examining these interactions without considering gender roles (e.g., Ngai and Pissarides, 2008) is a useful step; but given the significant differences in gender roles in less developed countries, understanding growth and development will require accounting better for the convergence of total work and changes in the relative amounts of market and household work performed by men and women. This is especially true if one considers the

different roles played by physical and intellectual attributes over the agricultural, industrial and service-sector phases of economic development (Clark, 1940; Fourastié, 1949).

Second, in household models we typically assume that a spouse's bargaining power is a function of her/his market earnings. Yet we have shown here, at least for most rich economies, that gender differences in the amounts of non-work time are tiny. How can this be true if, as is still the case, men have substantially higher wage rates and market earnings? Three logical possibilities present themselves. Men have more power, but are altruistic toward their spouses and toward women generally, and do not take advantage of it.²³ Another possibility is that economists' modeling of the household has been incorrect, and market earnings do not generate power in the household. A final alternative is that earnings do generate power, men are not altruistic, but the average man's utility from his market and home work exceeds that of the average woman's from the same total amount of work. In other words, iso-work may not imply iso-utility from the same amount of work. This last possibility would formalize ideas of the few sociologists who have confronted the iso-work phenomenon (e.g., Mattingly and Bianchi, 2003). Yet this possibility shifts the discussion to why women find their work more onerous than men find theirs, which, as it involves interpersonal utility comparisons, is not a question for economists. Why, e.g., is the marginal minute spent in an office dealing with recalcitrant colleagues and demanding supervisors more pleasurable than the marginal minute spent baking a cake or reading to a three-year-old?

Third, the choice between home and market work is generally determined by equating the marginal value product of household time to the real wage measured in terms of comparable market output (e.g., Gronau, 1980). To the extent that social norms constrain agents to a value of labor supply that ignores individuals' productivity in household production, one or more

²³Doepke and Tertilt (2008) present a model in which self-interest motivated by inter-generational concerns leads men to use their power to grant equal rights to women. Bertocchi (2007) constructs a model in which concerns by a changing median voter lead to extensions of rights to women.

efficiency conditions in the standard model of household behavior might be violated, or needs to be modified.

Our results indicate the potential importance of going beyond standard neoclassical models to analyze social phenomena. Iso-work does not prove the existence of social norms; but our inability to predict its patterns and correlates using neoclassical models and its consistency with the implications of social norms suggest that this is a fruitful general approach to the study of labor markets generally. It shows that an inductive, data-driven approach can be a useful guide to the development of theoretical discussions of the social determinants of economic relations.

REFERENCES

- Mark Aguiar and Erik Hurst, "Measuring Trends in Leisure: The Allocation of Time over Five Decades," *Quarterly Journal of Economics*, 122 (August 2007): 969-1006.
- George Akerlof, "A Theory of Social Custom, of Which Unemployment May Be One Consequence," *Quarterly Journal of Economics*, 94 (June 1980): 749-775.
- Christel Aliaga and Karin Winqvist, "How Women and Men Spend Their Time" *Statistics in Focus: Population and Social Conditions*. Eurostat, 2003.
- B. Douglas Bernheim, "The Economics of Conformity," *Journal of Political Economy*, 102 (October 1994): 841-877.
- Graziella Bertocchi, "The Enfranchisement of Women and the Welfare State," Unpublished paper, Università di Modena e Reggio Emilia, May 2008.
- Jeff Biddle and Daniel Hamermesh, "Sleep and the Allocation of Time," *Journal of Political Economy*, 98 (October 1990): 922-43.
- Michael Bittman and Judy Wacjman, "The Rush Hour: The Character of Leisure Time and Gender Equity," *Social Forces*, 79 (September 2000): 165-89.
- C. Mark Blackden and Quentin Wooden, eds., *Gender, Time Use, and Poverty in Sub-Saharan Africa*. Washington, DC: World Bank, 2006.
- George Borjas, "Ethnic Capital and Intergenerational Mobility," *Quarterly Journal of Economics*, 107 (February 1992): 123-150.
- Edward Cartwright, "On the Emergence of Social Conformity," Department of Economics, University of Kent, Discussion Paper No. 0501, 2005.
- Colin Clark, *The Conditions of Economic Progress*. London: Macmillan, 1940.
- Harold Cole, George Mailath and Andrew Postlewaite, "Social Norms, Savings Behavior and Growth," *Journal of Political Economy*, 100 (December 1992): 1092-1125.
- Matthias Doepke and Michèle Tertilt, "Women's Liberation: What's in It for Men?" Unpublished paper, Stanford University, March 2008.
- Jon Elster, "Social Norms and Economic Theory," *Journal of Economic Theory*, 3 (Fall 1989): 99-117.
- Jean Ensminger, "Market Integration and Fairness: Evidence from Ultimatum, Dictator, and Public Goods Experiments in East Africa," in Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr and Herbert Gintis, eds., *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence From Fifteen Small-Scale Societies*. Oxford: Oxford University Press, 2004, 356-81.

- Raquel Fernandez, "Women, Work and Culture," *Journal of the European Economic Association*, 5 (2007): 305-333.
- Jean Fourastié, *Le Grand Espoir du XXe Siècle; Progrès Technique, Progrès Économique, Progrès Social*. Paris: Presses Universitaires de France, 1949.
- Richard Freeman and Ronald Schettkat, "Marketization of Household Production and the EU-US Gap in Work," *Economic Policy*, 20 (2005): 6-50.
- Jonathan Gershuny, *Changing Times: Work and Leisure in Postindustrial Society*. Oxford: Oxford University Press, 2000.
- Claudia Goldin, "The U-Shaped Female Labor Force Participation Function in Economic Development and Economic History," in T. Paul Schultz, ed., *Investment in Women's Human Capital and Economic Development*. Chicago: University of Chicago Press, 1995, pp. 61-90.
- Luisella Goldschmidt-Clermont and Elisabetta Pagnossin-Aligisakis, "Measures of Unrecorded Economic Activities in Fourteen Countries," in *Occasional Paper # 20*. New York: UN Development Report Office, 1995.
- Jeremy Greenwood, Richard Rogerson and Randall Wright, "Household Production in Real Business Cycle Theory," in Thomas Cooley, ed., *Frontiers of Business Cycle Research*. Princeton, NJ: Princeton University Press, 1995, pp. 157-174.
- Jeremy Greenwod, Ananth Seshadri and Mehmet Yorukoglu, "Engines of Liberation," *Review of Economic Studies*, 72 (January 2005): 109-133.
- Reuben Gronau, "Home Production: A Forgotten Industry," *Review of Economics and Statistics*, 62 (August 1980): 408-416.
- Lawrence Haddad, Lynn Brown, Andrea Richter and Lisa Smith, "The Gender Dimensions of Economic Adjustment Policies: Potential Interactions and Evidence To Date," *World Development*, 23 (1995): 881-896.
- Daniel Hamermesh, "Timing, Togetherness and Time Windfalls," *Journal of Population Economics*, 15 (November 2002): 601-623.
- Alan Heston, Robert Summers and Bettina Aten, *Penn World Tables Versions 6.1, 6.2*. Philadelphia: Center for International Comparisons of the University of Pennsylvania (CICUP), 2002, 2006.
- Arlie Russell Hochschild, *The Time Bind*. New York, Henry Holt, 1997.
- Stephen R. G. Jones, *The Economics of Conformism*. Oxford: Blackwell, 1984.
- Michihiro Kandori, "Social Norms and Community Enforcement," *Review of Economic Studies*, 59 (January 1992): 62-80.
- Peter Kooreman and Arie Kapteyn, "A Disaggregated Analysis of the Allocation of Time within the Household," *Journal of Political Economy*, 95 (April 1987): 223-49.

- Assar Lindbeck, "Incentives and Social Norms in Household Behavior," American Economic Association, *Papers and Proceedings*, 87 (May 1997): 370-377.
- , Sten Nyberg and Jörgen Weibull, "Social Norms and Economic Incentives in the Welfare State," *Quarterly Journal of Economics*, 114 (February 1999): 1-35.
- Shelly Lundberg and Robert Pollak, "Bargaining and Distribution in Marriage," *Journal of Economic Perspectives*, 10 (Fall 1996): 139-158.
- Charles Manski, *Identification Problems in the Social Sciences*. Cambridge, MA: Harvard University Press, 1995.
- Marybeth Mattingly and Suzanne Bianchi, "Gender Differences in the Quantity and Quality of Free Time," *Social Forces*, 82 (March 2003): 999-1030.
- L. Rachel Ngai and Christopher Pissarides, "Trends in Hours and Economic Growth," *Review of Economic Dynamics*, 11 (April 2008): 239-256.
- Solomon Polachek and Jun Xiang, "The Gender Pay Gap: A Cross-Country Analysis," Unpublished paper, SUNY-Binghamton, February 2006.
- Margaret Reid, *Economics of Household Production*. New York: Wiley, 1934.
- John Robinson and Geoffrey Godbey, *Time for Life: The Surprising Ways Americans Use Their Time*, 2nd ed. University Park, PA: Pennsylvania State University Press, 1999.
- Pitirim Sorokin and Clarence Berger, *Time-Budgets of Human Behavior*. Cambridge: Harvard University Press, 1939.
- Yoram Weiss, "Work and Leisure: A History of Ideas," *Journal of Labor Economics*, 27 (2009): forthcoming.
- Myrna Wooders, Edward Cartwright and Reinhard Selten, "Behavioral Conformity in Games with Many Players," *Games and Economic Behavior*, 57 (November 2006): 347-360.
- H. Peyton Young, "The Economics of Convention," *Journal of Economic Perspectives*, 10 (Spring 1996): 105-122.

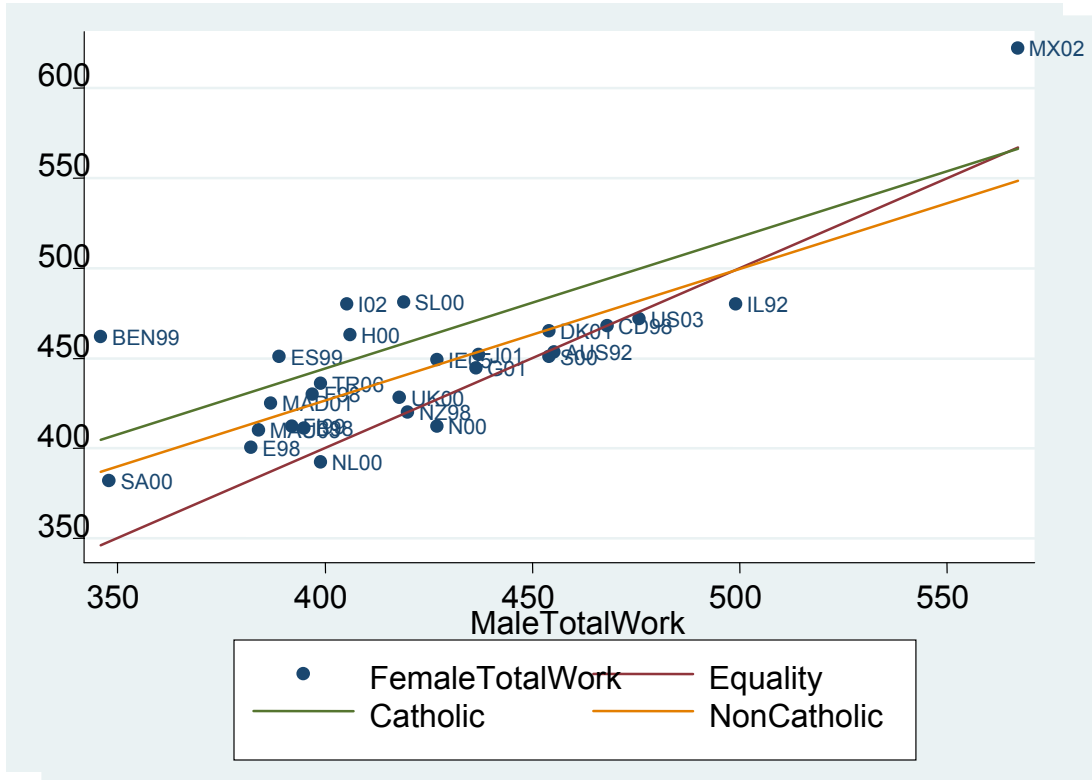


Figure 1. Scatter and Linear Regression of Female Total Work Against Male Total Work and Non-Catholic (Upper Parallel Line), Catholic (Lower Parallel Line), Equality of Total Work, 27 Countries*

*The country abbreviations are listed in the Appendix.

Table 1. Expert and Other Opinion about Men’s and Women’s Total Work, Percent Distributions and Statistical Tests

	Labor economists familiar with US	Labor economists unfamiliar with US	Elite macro and public finance economists	Sociology faculty and graduate students	Economics principles students
Men work:					
25% less	5.2	5.6	2.6	20.0	6.1
15% less	18.8	20.4	23.7	26.7	18.9
10% less	17.8	24.6	18.4	13.3	18.7
5% less	11.7	11.3	10.5	11.7	12.8
Differ by less than 2.5%	25.8	25.4	34.2	20.0	23.1
5% more	6.1	4.9	3.9	1.7	9.2
10% more	8.5	4.2	5.3	3.3	7.0
15% more	5.2	2.1	0	3.3	3.4
25% more	0.9	1.4	1.3	0	0.9
N =	213	142	76	60	445
Fraction with men < women	0.535	0.620	0.553	0.717	0.564
t-statistic on binomial if “equal” answers are split evenly	5.47	6.73	4.08	6.08	8.98
t-statistic on binomial if “equal” answers are assumed men > women	1.03	2.93	0.92	3.69	2.72
RESPONSE RATE	0.535		0.298	0.286	0.873

Responses are to the question: “We know that American men (ages 20-75) on average work more in the market than do American women. But what is the difference between men’s TOTAL WORK (in the market and on anything that you might view as work at home) and that of women?”

Without consulting any books, articles or raw data, PLEASE PUT AN X NEXT TO THE LINE BELOW THAT YOU BELIEVE TO BE THE CLOSEST APPROXIMATION TO THE CURRENT SITUATION IN THE US.”

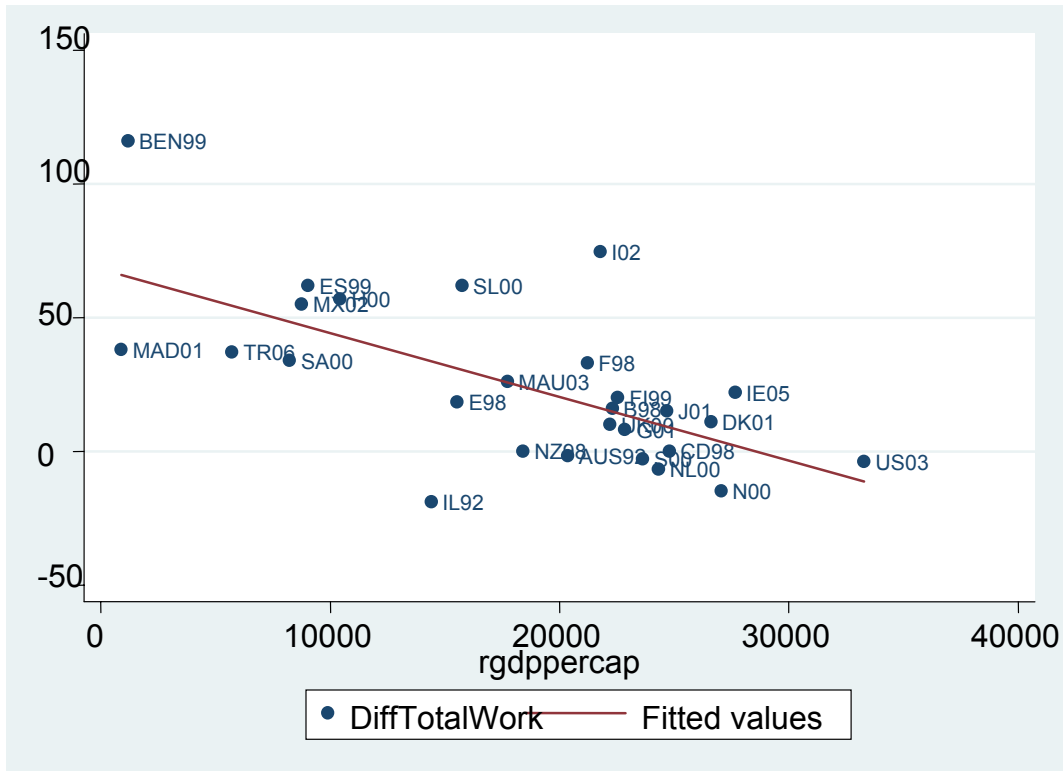


Figure 2. Difference between Female and Male Total Work Compared to Real GDP per Capita, 27 Countries

$$\text{DiffTotalWork} = 68.16 - 2.39\text{RealGDP/Capita}; R\text{Bar}^2 = 0.394$$

(0.56)

Table 2. Impact of the Gender Pay Gap on the Gender Total Work (minutes per day), N = 19*

Dep. Var., Female – Male Total Work:	(1) Market work	(2) Total work	(3) Market work	(4) Total work	(5) Market work	(6) Total work
Log (Female/Male Wage)**	135.5 (73.02)	44.0 (38.54)	182.8 (66.96)	23.47 (37.42)	162.68 (71.71)	55.00 (32.34)
Real GDP per capita (\$000)***			4.38 (1.78)	-1.90 (1.00)	3.98 (1.92)	-1.20 (0.87)
Catholic					-17.22 (26.07)	33.78 (11.76)
Adj. R ²	0.120	0.017	0.321	0.149	0.293	0.423

*Standard errors in parentheses here and in Tables 6 and 7.

**From Polachek and Xiang (2006).

***From Heston *et al* (2006).

Table 3. Time Allocations (minutes per representative day), Averages and Their Standard Errors, Women, Married and Unmarried Separately, U.S. 2003, Germany 2001/02*

	U.S. 2003, Married		U.S. 2003, Unmarried		Germany, 2001/02, Married		Germany, 2001/02, Unmarried	
	F	M	F	M	F	M	F	M
Market work	182 (3.4)	329 (4.4)	224 (4.1)	284 (5.5)	111 (2.1)	270 (3.3)	175 (3.8)	241 (5.5)
Home work	314 (2.8)	179 (2.6)	218 (2.8)	136 (3.0)	336 (2.0)	175 (1.8)	264 (2.9)	170 (3.5)
Female- Male Total work		-12		22		2		28

*Standard errors of means here and in Tables 4 and 5.




Figure 3a. Histogram of Wife-Husband Differences in Total Work, United States 1985




Figure 3b. Histogram of Wife-Husband Differences in Average Total Work per Day over Two Days, Australia 1992)




Figure 3c. Histogram of Wife-Husband Differences in Average Total Work per Day over Three Days, Germany 2001/02

Table 4. Total Work by Education Level, U.S. 2003, Germany 2001/02

	Highest	2nd	3rd	Lowest
Gender	United States			
F	518 (4.33)	474 (4.4)	455 (4.5)	386 (7.6)
M	524 (5.1)	470 (6.1)	468 (5.9)	366 (9.2)
	Germany			
F	172 (2.8)	152 (6.5)	147 (2.9)	98 (3.0)
M	270 (4.8)	290 (8.2)	273 (4.9)	237 (4.4)
F	475 (3.6)	465 (6.2)	456 (2.8)	406 (3.3)
M	455 (4.2)	456 (7.2)	448 (4.2)	416 (4.0)

Table 5. Total Work by Region, Ages 20-74, U.S. 2003, Germany 2001/02

Gender	United States	
	Non-South	South
F	480 (3.0)	457 (4.2)
M	480 (3.8)	467 (5.5)
	Germany	
	West	East
F	445 (2.0)	465 (4.0)
M	436 (2.5)	451 (5.2)

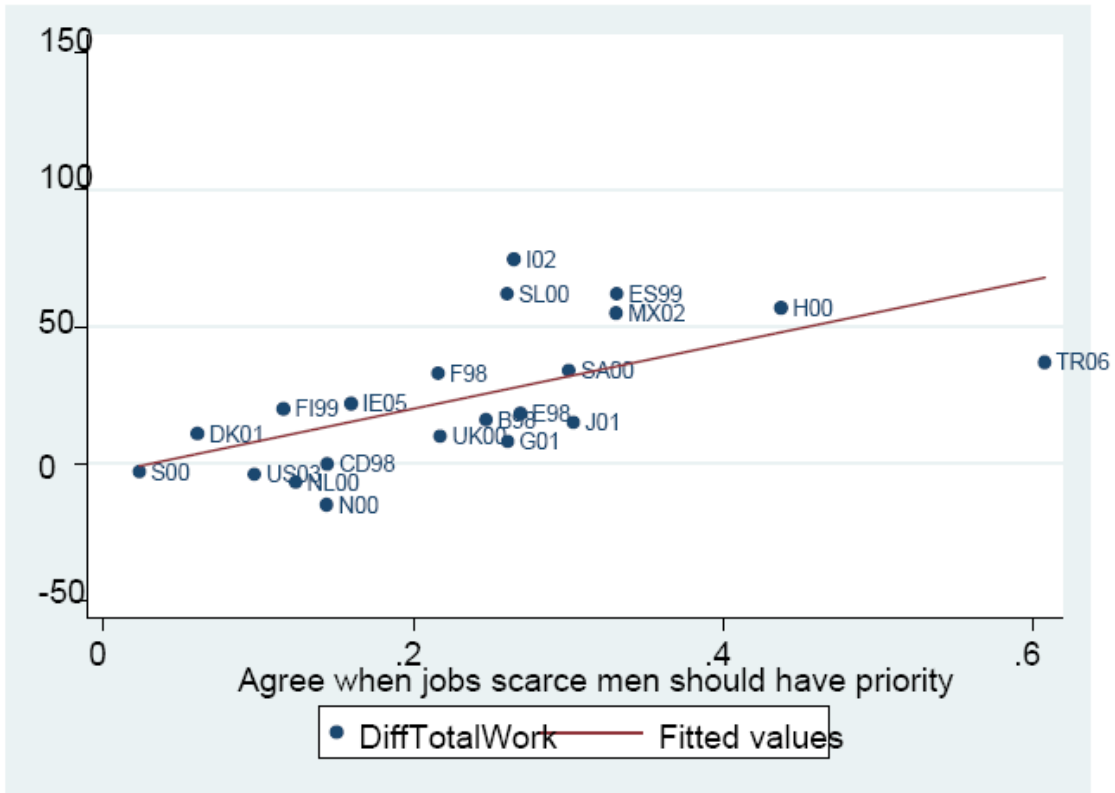


Figure 4a. Gender Total Work Differences and Average Attitudes to Job Scarcity, 21 Countries

$$\text{DiffTotalWork} = -3.39 + 117.50\text{AgreeJobsMen}; R\text{Bar}^2 = 0.337$$

(35.16)

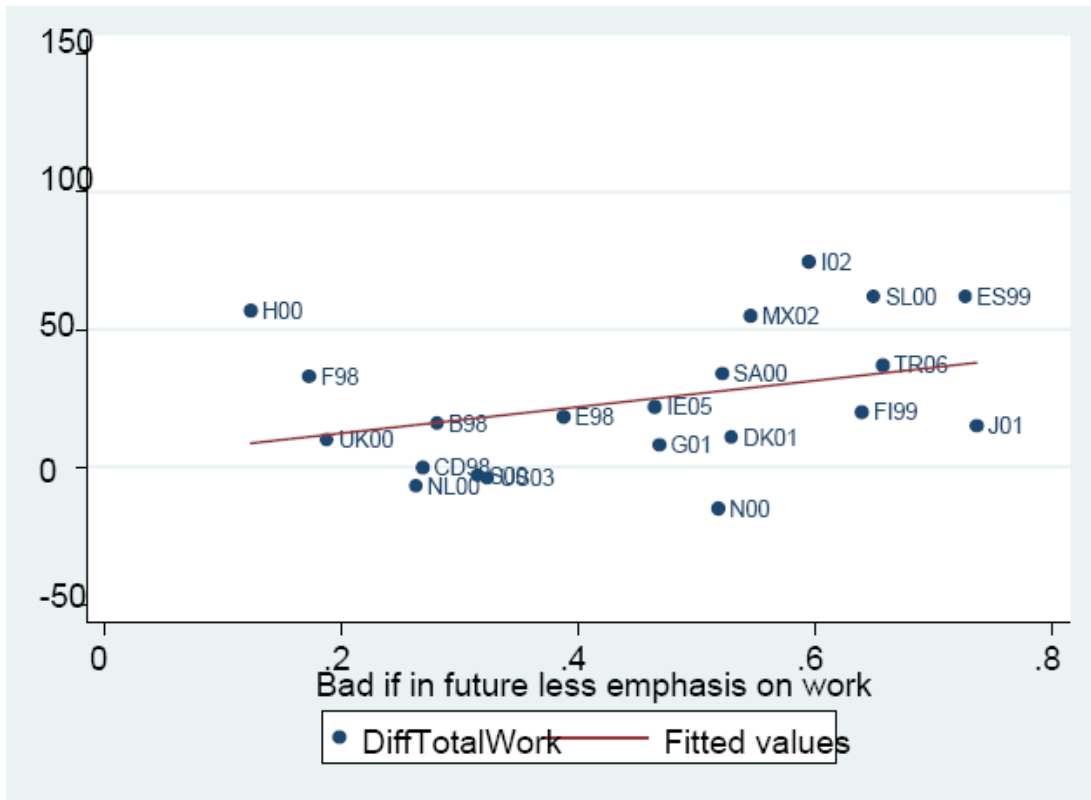


Figure 4b. Gender Total Work Differences and Average Attitudes to Value of Market Work, 21 Countries

$$\text{DiffTotalWork} = 2.87 + 47.67\text{LessWorkBad} ; R\text{Bar}^2 = 0.078$$

(29.32)

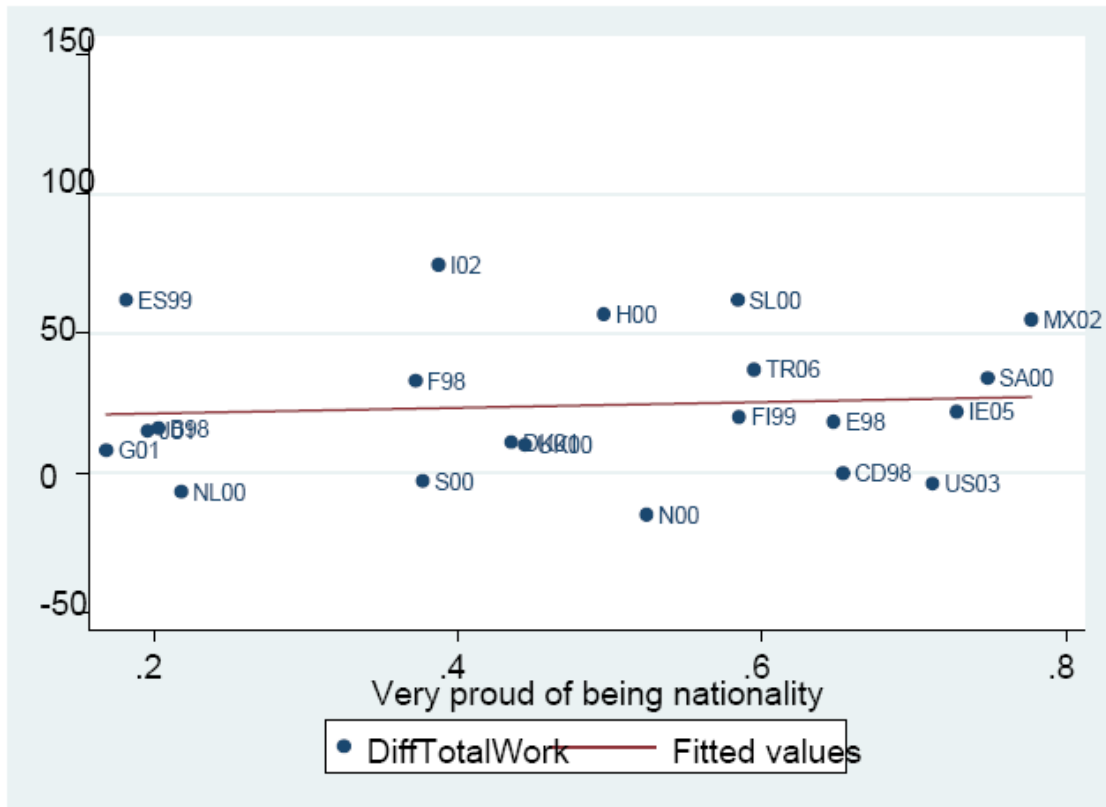


Figure 4c. Gender Total Work Differences and Average Pride in Nationality, 21 Countries

$$\text{DiffTotalWork} = 19.24 + 10.22\text{VeryProudNationality}; R\text{Bar}^2 = -0.046$$

(29.11)

Table 6. The Effects of Home Country Social Norms on Immigrants and Second-Generation Americans Total Work Time, 2003-2006 (Dep. Var. is Wife's Total Work Time)^a

	Immigrants			Second Generation		
	(1)	(2)	(3)	(4)	(5)	(6)
Agree Men Should Have Job Priority	105.84 (48.87)			-45.16 (131.24)		
Bad If Less Emphasis On Work		-52.72 (45.30)			-36.20 (120.83)	
Very Proud of Nationality			21.12 (34.54)			-59.90 (98.70)
Total Work Husband	-0.037 (0.251)	0.212 (0.248)	0.184 (0.252)	0.119 (0.223)	0.100 (0.218)	0.102 (0.218)
RBar ²	0.095	-0.040	-0.088	-0.144	-0.145	-0.130
N =		26			25	

^aAlso included in the regressions are the average ages of women and men in the group.

Table 7. Direct Tests for the Effects of Social Norms, Married Couples Germany 2001/02, Australia 1992 (Dep. Var. is Wife's Total Work Time)^a

	Germany				Australia		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Total Work Husband	0.491 (0.012)	0.441 (0.013)	0.432 (0.013)	0.447 (0.013)	0.416 (0.015)	0.358 (0.016)	0.355 (0.016)
Average Total Work in Wife's: Education Group		0.298 (0.081)	0.252 (0.075)	-0.404 (0.231)		0.466 (0.111)	0.587 (0.129)
Age Group		0.281 (0.046)	0.055 (0.065)	0.216 (0.033)		0.431 (0.056)	0.231 (0.070)
Region		0.722 (0.159)	0.873 (0.161)	0.443 (0.098)		1.276 (0.518)	1.668 (0.812)
Norm Used		Current	Current	Lagged		Current	Current
Wife's Education, Age			Yes				Yes
Family's Number and Ages of Children			Yes				Yes
RBar ²	0.353	0.372	0.377	0.368	0.292	0.325	0.336
N=		3080				1966	

^aIn Germany total work is the average over three days, in Australia it is the average over two days. For both countries the age groups are under 40, 40-54 and 55+. There are four education groups in Germany, three in Australia. Germany is divided into West and East, Australia in New South Wales, Victoria, Queensland and other jurisdictions.

DATA APPENDIX: Definitions of Total Work in 27 Countries

United States(US): Market work and work-related activities; travel related to work; all household activities; caring for and helping household members; consumer purchases; professional and personal care services; household services; government services; travel related to these.

Australia(AUS): Market work; cleaning and cooking; family and child care; shopping; and travel associated with each.

Belgium(B), Denmark(DK), France(F), Finland(FI), Sweden(S), United Kingdom(UK), Estonia(ES), Hungary(H), Slovenia(SL) Norway(N): Gainful work; study; household work + family care; proratio of travel time based on gainful work time.

Benin(BEN), Madagascar(MAD), Mauritius(MAU), South Africa(SA): Market work; domestic and care activities; commuting.

Canada(CD): “Total work” (paid work and related activities; unpaid work and related activities).

Germany(G): Market work: employment and job search; home work activities; handicraft/gardening; care and sitting.

Ireland(IE): Care; employment and study; household work; proratio of travel time based on gainful work time.

Israel(IL): Market work; cooking and cleaning at home; child care.

Italy(I): market work; professional activities; training; domestic activities; family care; purchasing goods and services.

Japan(J): Work, school work; house work, caring or nursing, child care, shopping.

Mexico(MX): Domestic work; care of children and other household members; market work.

Netherlands(NL): Occupational work and related travel; household work, do-it yourself, gardening, etc; childcare; shopping.

New Zealand(NZ): Paid work; household work, care-giving for household members, purchasing goods or services, unpaid work for people outside the home.

Spain(ES): Market work; house work, child care, adult care.

Turkey(TR): Employment and job seeking; study; household and family care; proratio of travel time based on gainful work time.

Appendix. Strategic complementarities: A Model of Stakhanov

In the text we have assumed that the cost of deviating does not depend on how many people conform. Plausibly, the stigma attached to deviating from a social norm (or the very existence of a social norm) depends on how widely accepted the norm is. For instance, the productivity norm established by Alexei Stakhanov (1906-1977), the legendary Soviet coal miner who in 1935 extracted fourteen times his quota, was compelling to individuals because his example was emulated, under the pressure of Soviet propaganda, by a large number of workers.

This consideration opens the possibility of multiple equilibria through the existence of strategic complementarities: if the loss that we experience when we deviate from the norm depends on how widely the norm is followed, whether or not we choose to conform depends on our perception of the prevalence of the norm among our fellow citizens. If we expect them to conform, we have an incentive to act as they do, for the cost of deviating is then high. If we anticipate that others will disregard the norm, however, the cost of deviating is small, and it is more likely that we will find it optimal to follow our intrinsic optimum. The question then arises whether deviating is the stable outcome that will emerge endogenously from a population of self-interested individuals. The answer, as we establish now, depends crucially on the shape of the wage distribution.²⁴

For simplicity, let us depart from the linear-quadratic model in the text, and assume instead that the cost of deviating is *fixed* from the point of view of individuals. However, let us imagine that it depends on the fraction $\pi \in (0, 1)$ of “conformists” in the population:

$$\psi = \pi^2/2. \tag{26}$$

If no one in society conforms, there is no cost of deviating. The cost of deviating from the norm is increasing in π , and thus maximal when everyone else conforms ($\pi = 1$).²⁵ One can easily show that an individual with wage w conforms if and only if

$$\psi = \frac{1}{2}\pi^2 > \frac{1}{2}(w - w^*)^2, \tag{27}$$

²⁴This Appendix is inspired by the work of Cartwright (2005) and Wooders *et al.* (2006) on the emergence of social conformity.

²⁵These two properties are crucial. The quadratic specification for ψ is adopted for simplicity but it is not innocuous, as the number and stability of equilibria depend jointly, as we will see below, on the shape of the cost function and on the distribution of wages.

that is, if and only if his/her wage is in the band $[w^* - \pi, w^* + \pi]$. The more widely adopted the norm is, the wider is the conformity band, and the more likely it is that an individual with an arbitrary wage will conform. Conversely, when fewer people conform, the narrower is the band, and the more likely it is that an individual will deviate. This strategic complementarity opens the door to multiple equilibria.

To illustrate this point, suppose male and females have the same cumulative wage distribution $F(w)$ over the interval $[0, 1]$. The fraction of the population with wages in the conformity band is then $F(w^* + \pi) - F(w^* - \pi)$. Since this fraction must coincide with π in equilibrium, the fraction of conformists in the population solves the equation

$$F(w^* + \pi) - F(w^* - \pi) = \pi. \quad (28)$$

Regardless of the exact shape of $F(\cdot)$, this equation always has at least two solutions, $\pi = 0$ and $\pi = 1$. The former corresponds to a *non-conformist* equilibrium in which no one conforms, and a *conformist equilibrium* in which everybody adheres to the norm.²⁶

Which of these equilibria is stable depends on the shape of the cumulative distribution function $F(\cdot)$. To illustrate this point, assume wages are distributed uniformly over $[0, 1]$, so that $F(z) = z$ for $0 < z < 1$, and $F(z) = 1$ for $z > 1$. In addition, assume for the moment that the norm is the median value of intrinsic leisure, so that $w^* = 1/2$ (i.e., the median of the individual w 's). Then

$$F(w^* + \pi) - F(w^* - \pi) = \begin{cases} 2\pi, & \text{for } 0 \leq \pi \leq 1/2; \\ 1, & \text{for } 1/2 < \pi \leq 1. \end{cases} \quad (29)$$

Figure A1 shows that for a uniform distribution of wages there are exactly two equilibria, $\pi = 0$ and $\pi = 1$. Crucially, only the conformist equilibrium is stable as $F(w^* + \pi) - F(w^* - \pi) > \pi$ for all π strictly between 0 and 1.²⁷ Hence full conformity to $L^* = 1 - w^* = 1/2$ will emerge endogenously in this economy, and the iso-work fact will hold in its strictest form even though agents have different wages and the cost of deviating from the norm is finite.

Remarkably, this reasoning holds regardless of the value of the norm. Suppose w^* is different

²⁶If $\pi = 1$, $w^* + \pi > 1$ so that $F(w^* + \pi) = 1$, while $w^* - \pi < 0$ so that $F(w^* - \pi) = 0$.

²⁷The easiest way to see this is to observe that the difference equation $F(w^* + \pi_t) - F(w^* - \pi_t) = \pi_{t+1}$ converges to $\pi = 1$ as $t \rightarrow \infty$ for any $0 < \pi_0 < 1$.




Figure A1: Multiple equilibria ($w^* = 1/2$; uniform distribution of wages over $[0, 1]$)

from $1/2$. Then, if we maintain the assumption that the distribution of male and female wages is identical and uniform over the interval $[0, 1]$, one can show that

$$\begin{aligned}
 & F(w^* + \pi) - F(w^* - \pi) \\
 &= \begin{cases} 2\pi, & \text{for } 0 \leq \pi \leq \min(w^*, 1 - w^*); \\ w^* + \pi, & \text{for } \min(w^*, 1 - w^*) < \pi \leq \max(w^*, 1 - w^*); \\ 1, & \text{for } \max(w^*, 1 - w^*) < \pi \leq 1. \end{cases} \quad (30)
 \end{aligned}$$

As before, there are two equilibria, illustrated in Figure A2: a stable one in which everybody conforms, and an unstable one in which nobody conforms. Nothing pins down the norm: w^* , and thus L^* , can take *any* value in the interval $[0, 1]$. Hence there is a continuum of equilibria with full conformity over $[0, 1]$, indexed by the social norm L^* .²⁸

What can we say about the welfare properties of these conformist equilibria? Can they be Pareto-ranked? To answer that question, we need only look at the welfare of an agent with wage w in the conformist equilibrium indexed by w^* , and examine how it depends on w^* . We established earlier that

$$U^C = \Omega + ww^* - \frac{1}{2}w^{*2}, \quad (31)$$

²⁸Note, however, that for other distributions (for instance, distributions with mass concentrated on extreme values), the non-conformist equilibrium might emerge as the stable one. In addition, one can construct examples in which $\pi = 0$ and $\pi = 1$ are not the only possible solutions, and in which the equilibrium fraction of conformists is strictly between zero and one, and stable. We do not explore these refinements here, but they might help us explain why some groups or countries experience less conformism than others.




Figure A2: Multiple equilibria (w^* arbitrary; uniform distribution of wages over $[0, 1]$)

so that

$$\frac{\partial U^C}{\partial w^*} = w - w^*. \quad (32)$$

Since $L^* = 1 - w^*$, this implies that

$$\frac{\partial U^C}{\partial L^*} = w^* - w. \quad (33)$$

Hence low-wage agents ($w^* - w > 0$) are better off in an economy in which social pressure dictates high leisure. Conversely, high-wage individuals ($w^* - w < 0$) are better off in a “stakhanovist” society in which L^* is high. This difference in welfare stems only, in our model, from the fact that agents prefer norms that are congruent with their intrinsic tastes. Hence, the continuum of conformist equilibria cannot be Pareto-ranked.