

Worktime Regulations and Spousal Labor Supply[†]

By DOMINIQUE GOUX, ERIC MAURIN, AND BARBARA PETRONGOLO*

We study interdependencies in spousal labor supply by exploiting the design of the French workweek reduction, which introduced exogenous variation in one's spouse's labor supply, at constant earnings. Treated employees work on average two hours less per week. Husbands of treated women respond by reducing their labor supply by about half an hour, consistent with substantial leisure complementarity, and specifically cut the unusual component of their workweek, leaving usual hours unchanged. Women's response to their husband's treatment is instead weak and rarely statistically significant, possibly due to heavier constraints in the organization of their workweek. (JEL J16, J22, K31)

Interdependencies in spousal labor supply have long been identified as a key question in the study of household behavior (Ashenfelter and Heckman 1974). Complementarities in labor supply and leisure within or beyond the household are also a key policy issue, as they represent a channel through which reforms targeted at specific segments of the population can ultimately affect a wider set of individuals. When the value of leisure time for an individual depends on the amount of leisure enjoyed by her spouse, coworkers, neighbors, social contacts, etc., reforms of the welfare state, or tax reforms, or changes in workweek regulations aimed at some segments of the workforce may impact individual behavior well beyond the targeted population (Alesina, Glaeser, and Sacerdote 2006).

While interdependencies in work and leisure represent an important and controversial issue, there is still little micro-level evidence on the actual magnitude of these effects. Progress in this direction has been limited by the difficulty of finding independent variation in the labor supply of one's peers, as individuals within the same family or social network may be subject to the same reforms, or more in general to correlated labor supply shocks. Another major challenge is that changes in leisure time and working hours are in most cases associated with important changes in earnings. Thus the labor supply responses of peers cannot be interpreted as reflecting

* Goux: CREST, 60 rue Etienne Dolet 92240 Malakoff, France (e-mail: dominique.goux@ensae.fr); Maurin: Paris School of Economics, 48 Boulevard Jourdan, 75014 Paris, France (e-mail: eric.maurin@ens.fr); Petrongolo: Queen Mary University, Mile End Road, London E1 4NS, UK, and CEP, LSE, Houghton Street, London WC2A 2AE, UK. (e-mail: b.petrongolo@lse.ac.uk). We thank Philippe Askenazy, François Bourguignon, Shelly Lundberg, and Thomas Piketty for helpful discussions as well as seminar participants at London School of Economics, Paris School of Economics, CREST Malinvaud Seminar, Bocconi University, University of Bologna, UC3M (Madrid), CEMFI (Madrid), ISER-Essex, University of Maastricht, Greqam (University of Aix-Marseille), the OECD, the CEPR conference on the Economics of Family Policy (Bergen 2010), and the CHILD Conference (Modena 2012) for valuable comments. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†] Go to <http://dx.doi.org/10.1257/aer.104.1.252> to visit the article page for additional materials and author disclosure statement(s).

pure cross-hour effects, as they may also encompass income effects. In this paper we exploit the unique design of the workweek reduction policy implemented in France in the late 1990s to overcome these issues and provide one of the very first micro estimates of the effect of an exogenous change in individuals' working hours on the labor supply of their spouses.

In June 1998 the French socialist government mandated a reduction of the legal workweek, from 39 to 35 hours, to be implemented at constant monthly earnings. This made the legal workweek in France (by far) the shortest among OECD countries (Lee, McCann, and Messenger 2007, Table 2.4). In order to attenuate the impact of higher hourly wages on profitability, firms who would implement the shorter workweek before the relevant deadline would benefit from significant payroll tax cuts. Only about 300,000 firms had implemented the shorter workweek before the comeback of the conservative party to power in April 2002 and the interruption of the original reform. Nevertheless, the reform implied a noticeable change in the workweek of at least one spouse in over one-third of French households, with no direct impact on family income. Both within-household variation in the workweek reduction, and the absence of income effects, make the French work-sharing reform a unique scenario for assessing cross-hour effects within the household.

In general, it is theoretically ambiguous whether a fall in working hours and thus an increase in non-market time of one spouse would generate a fall or a rise in working hours of the other spouse. Substitutability in non-market time of husbands and wives could be driven by substitutable spouse efforts in home production. A reduction in the workweek of one spouse may shift some of her time endowment from market to home production, thus freeing-up some home production time of the other spouse, who could devote more time to market work. Conversely, if one detects complementarity in the non-market time of spouses, this would rather be consistent with complementarity of their leisure time. A reduction in the workweek of one spouse would increase her leisure time and thus raise the value of leisure of the other spouse if spouses enjoy spending time together.

This paper uses a matched worker-firm dataset obtained by combining the French Labor Force Survey with firm-level information on the implementation of the shorter workweek, in order to estimate the labor supply response of men and women to a reduction in the legal workweek in their spouses' workplaces. We detect an average reduction of about 2 hours in the workweek of employees whose employers signed a workweek reduction agreement.¹ When looking at spousal responses, we find that men tend to work about half an hour less per week when their wives become treated, while women's response to their husbands' treatment is generally weak and rarely significantly different from zero.

Further tests reveal that men's labor supply response to wife treatment is not associated with a reduction in their usual working hours, but with a reduction in the "nonusual" component of their workweek. Moreover, such response does not have a detrimental impact on their earnings, suggesting that men manage to cut on some form of unpaid work involvement, whether within a given day, or through an increase in the take-up rate of paid vacation and/or sick leave. If employees do not use their

¹ We will discuss below various reasons why the average effect of the shorter legal workweek on actual weekly hours is lower than the legal workweek reduction.

whole paid leave entitlement, or do some unpaid overtime, they have some leeway in cutting their hours while avoiding earnings losses, and it is mostly by adjusting around these unpaid work margins that men respond to shorter workweek agreements in their wives' firms. Under the assumption that the workweek reduction in wives' firms affects their husbands only via wives' labor supply, we provide an instrumental variable estimate of the average cross-hour effect for husbands of 0.23, rising to 0.34 for managers and professionals, and to 0.59 for fathers of young children.

Our paper builds on a long strand of literature on family labor supply, investigating the response of an individual's labor supply to independent changes in her spouse's income and/or hours of work. These changes may in turn be driven by events as diverse as retirement, job loss, or fiscal reforms. Several studies document the positive association between husbands' and wives' retirement decisions, over and above what would be predicted by correlation in age and incentives in the retirement system (Blau 1998; Gustman and Steinmeier 2000). Conversely, the added worker effect literature detects mild substitutability between the labor supply of spouses, as married women tend to increase their working hours following husbands' job loss (Lundberg 1985; Cullen and Gruber 2000). More recently, Gelber (forthcoming) exploits the Swedish tax reform of 1990–1991 to examine own earnings' responses to changes in the marginal tax rate for one's spouse, and shows that as spousal earnings rise, own earnings rise too. Insofar as earnings responses reflect labor supply responses, these findings suggest complementarity in spousal leisure. Complementarity is also detected by Hamermesh (2002), who finds that spouses' daily work schedules are more synchronized than would occur randomly. While building on very different sources of variation, these papers agree in documenting important spillovers in the labor supply of spouses.

Our contribution to this literature is threefold. First, independent variation in spousal hours of work at constant earnings allows us to obtain cross-hour effects that are not confounded by income effects. In particular, under the assumption that an employee's workweek regulations affect their spouses only via their labor supply, we can identify the presence of leisure complementarity in the utility functions of spouses. Second, while previous work has mostly focused on the labor supply response of secondary earners, we find that it is in fact men who more strongly respond to their wives' treatment, while the corresponding women's response is much weaker. This may in turn be due to different degrees of leisure complementarities in the utility functions of spouses, or to a greater ability of men to control their working schedules. While we do not find compelling evidence on different preferences, the fact that women work shorter hours in the first place, and are less likely than men to hold managerial positions, suggests that they face relatively more binding constraints in the organization of their working time. Third, we provide evidence on specific adjustment margins in labor supply spillovers, and in particular we find that it is mostly men's unusual, rather than usual, hours that are affected when their wives' workweek is reduced.

In addition to the literature on household labor supply, our paper relates to another strand of the labor supply literature, investigating differences between micro and macro labor supply elasticities. Macroeconomic calibrations typically imply much larger labor supply elasticities than microeconomic estimates (Chetty et al. 2011a, b), and the recent literature has investigated two main channels potentially driving such a gap. First, work on social multipliers illustrates how social interactions would magnify aggregate responses relative to individual behavior in

a range of contexts, including labor supply (Glaeser, Sacerdote, and Scheinkman 2003; Maurin and Moschion 2009). Second, recent studies on optimization frictions have shown that costs of adjusting working hours at the intensive margin attenuate micro elasticities relative to aggregate responses (Chetty et al. 2011a; Chetty 2012).

Our work contributes to the understanding of mechanisms underlying either channel and the interaction between them. Specifically, labor supply spillovers are substantially shaped in nature and magnitude by optimization frictions, insofar as the cost of adjusting working hours restricts spousal labor supply responses to workers who have fewer constraints in organizing their workweek, and to the nonusual component of their workweek. The resulting labor supply spillovers are thus strongly asymmetric, whereby women's treatment affects male labor supply but not vice versa (with very few exceptions), and independent changes in usual working hours produce spillovers on nonusual hours, but not vice versa. Spillovers on nonusual hours may in turn have an impact on productivity and profitability, while the absence of spillovers on usual hours would in most cases rule out an impact on current earnings. As optimization frictions in working hours are likely to bind in a variety of institutional contexts, the French case study considered here can shed light on the nature and magnitude of labor supply spillovers in other scenarios.

Finally, our paper contributes to the literature on work-sharing policies in developed countries.² The study which is closest to ours is Hunt (1998), who shows that the gradual decline in standard working hours of male employees between 1985 and 1995 in Germany was not accompanied by changes in their wives' employment rates, but by a small decline in their hours of work. This result, while consistent with complementarity in spousal leisure, may also reflect wives' own gradual exposure to shorter standard workweeks.

The paper proceeds as follows. Section I gives an overview of the workweek reduction reform. Section II describes the data used and provides some graphic evidence. Section III presents our main regression results. Section IV addresses a number of caveats to a causal interpretation of our estimates. Section V provides instrumental variable estimates of cross-hour effects, using mandated workweek reductions as instruments for spousal labor supply. Section VI concludes.

I. Historical and Institutional Context

Since the early 1980s, the legal workweek duration in France has been 39 hours, accompanied by a 25 percent overtime wage premium and a 130 overtime hour limit per worker per year. This scenario was substantially changed in the late 1990s. In April 1997, the French president Jacques Chirac dissolved the parliament and called general elections one year ahead of the end of the legislature. This decision was highly unexpected and the electoral campaign that followed was very short. The socialist party proposed a program whose main axis was the reduction of unemployment through work-sharing, with two basic slogans: "*travailler moins pour travailler tous*" (work less in order to work all) and "*35 heures payées 39*" (35 worked hours paid 39). The left coalition won the election in June 1997.

²The employment effects of workweek reduction reforms in France are studied by Crépon and Kramarz (2002); Askenazy (2008); Esteveo and Sa (2008); and Chemin and Wasmer (2009), among others.

The workweek reduction was implemented in two steps (see Askenazy 2008, for a detailed description of the reform). The first law (*Aubry I*, after the then labor secretary Martine Aubry), passed in June 1998, set the legal workweek at 35 hours in the private sector and mandated its implementation by January 2000 in firms with more than 20 employees, and by January 2002 in smaller firms.³ Hours worked beyond the thirty-fifth hour would be treated as overtime hours. Firms who would implement the shorter workweek through collective agreements with unions before the relevant deadline would benefit from substantial cuts in payroll taxes,⁴ provided that they committed to maintain employment levels. Finally, the law required that workers should not experience a drop in their monthly earnings following the legal workweek reduction.⁵ In particular, firms who signed a 35-hours agreement had to grant a specific (4 hours) bonus to workers paid the monthly minimum wage. The general purpose of the law was to induce firms to raise employment levels by work-sharing, while offering them fiscal advantages to attenuate detrimental impacts of work-sharing on profitability.

In January 2000, the second law (*Aubry II*) introduced a few amendments in order to limit the burden of the shorter workweek on employers. Specifically, with a slight redefinition of working time, it made it possible for employers to exclude “unproductive breaks” from the hours count, and thus achieve some reduction in the measure of working hours without changing work schedules. Also, it allowed firms to implement shorter hours on an annual—rather than weekly—basis, with a 1,600 annual hour cap. This means that fiscal advantages could be obtained even with actual workweek reductions below 10 percent. Finally, the *Aubry II* law introduced a two-year transitional phase during which it was possible for employers to keep the 39-hour workweek by using overtime at a reduced 10 percent rate.

Two years later, in summer 2002, the conservative party came back to power and, while the *Aubry* laws remained formally in place, the transition to the shorter workweek was discontinued in practice. The new government raised the maximum number of overtime hours from 130 to 220, and extended fiscal incentives to all firms, including those that did not sign workweek reduction agreements. In this new scenario firms could effectively have employees working 39 hours weekly, at no extra hourly cost with respect to the pre-reform scenario. Following these political changes, the 35-hour workweek was never fully implemented, especially in small private firms. Nevertheless, the *Aubry* laws have had a very large impact on the French economy, covering about 10 million workers by 2002.

In a nutshell, the French workweek reform had several important features: it was largely unexpected; it has been interrupted, with only a fraction of workers being affected; it did not affect monthly earnings; and given its gradual implementation it would likely not treat spouses in a given household at the same time. We build on these features of the reform in order to evaluate the effect of an exogenous variation in an employee’s workweek on the labor supply of her spouse.

³There were no explicit deadlines set for firms in the public sector.

⁴For workers paid at the minimum wage, the cuts imply a reduction of about 8 percent in total labor cost for five years.

⁵As in principle there might be an income effect through overtime pay, we will illustrate in Section III the (lack of) earnings effects of the shorter workweek.

II. Data and Descriptive Evidence

A. The Dataset

We combine individual level information on worker characteristics and working hours with firm level information on collective agreements signed by employers who adopted the shorter workweek. Individual level information comes from the French Labor Force Survey (LFS), which is conducted by the French Statistical Office, INSEE. Before 2003, the LFS was conducted in March of every year, and covered a representative sample of about 100,000 households each year (with a 1/300 sampling rate). Since 2003, the survey is conducted each quarter and covers a representative sample of about 55,000 households each quarter. Our main analysis will be based on all repeated cross sections from 1994–2009, namely all annual surveys 1994–2002, and all first-quarter surveys for 2003–2009.

For each household member aged 15 or above, the LFS provides information on gender, marital status, employment status, occupation, education, industry, monthly earnings, and hours worked. We exploit information on both actual hours worked during the reference week (typically the week before the survey), and usual hours worked in a typical week.⁶ Crucial for our purposes, our restricted use version of the LFS also provides coded employer identifiers.⁷ These allow us to match worker level information with firm level information from the DARES-URSSAF dataset, an administrative database collected by the French Ministry of Labor, which provides detailed information on all firms who signed a workweek reduction agreement, including the signing and implementation dates. We thus obtain a matched employer-employee dataset containing information on working hours of respondents and their spouses, as well as information on when, if ever, their employers implemented the shorter workweek. The matched employer-employee dataset used has some clear advantages compared to the non-matched LFS. First, it allows us to identify which workers were actually treated, and not simply the intention to treat based on the number of employees in their firms and the proximity to the law deadlines. Also, the information on the exact date of treatment makes it possible to exploit the gradual implementation of the shorter workweek, thus avoiding to solely rely on the announced 2000 and 2002 deadlines. Detailed information on the construction of our dataset is provided in Goux, Maurin, and Petrongolo (2014).

In our analysis we select all married or cohabiting individuals aged 18–65, whose spouse is a wage earner, and we focus on the labor supply response of these individuals to their spouses' exposure to the shorter workweek. We define treatment as

⁶According to the official International Labour Office (2002, p. 203) definition, usual hours represent "the modal value of the number of hours actually worked per week over a long period of time." This definition is applicable to workers with regular schedules only (about 85 percent of cases in the LFS), and does not include irregular or unusual overtime, nor unusual absences or rest. French labor laws require contracts to be explicit about hours, pay, tasks, and paid leaves, and as a consequence interviewees would know precisely their normal hours as well as contractual changes in these. Moreover information in the LFS is collected through face-to-face interviews during which INSEE interviewers attempt to make sure that respondents understand questions and answer in a consistent way. This procedure considerably reduces measurement error on hours of work relative to self-filled questionnaires (Baum-Snow and Neal 2009).

⁷Each employee is asked to report the name and address of her employer, and this information is coded by INSEE. The coded employer identifier is available for just over 80 percent of the employees in the LFS. Most cases of missing employer ID correspond to very small firms. For a detailed description of the coding procedure, see Abowd and Kramarz (1999) or Goux and Maurin (1999).




FIGURE 1. TIMING OF IMPLEMENTATION OF THE SHORTER WORKWEEK:
PERCENTAGE OF EMPLOYEES TREATED

working for an employer who has signed a workweek reduction agreement,⁸ and we drop the small number of individuals whose spouses were treated either before 1996 or after 2002, since these early and late agreements may not correspond to the reform implemented in the late 1990s. Our working sample includes 189,894 males and 236,802 females. Descriptive statistics on these samples are provided in Table A1 and Figures A1 and A2 of the online Appendix.

To illustrate the timing of treatment, Figure 1 shows the gradual implementation of the shorter workweek on our sample. While only about 40 percent of employees are eventually treated, there is substantial variation in treatment dates between 1998 and 2002. Table 1 reports the distribution of own and spousal treatment for employed respondents, and shows that about 54 percent of husbands of treated women are not treated themselves by the workweek reduction (panel A, column 2), while about 29 percent of husbands of non-treated women are treated. While there is some assortative mating along the treatment dimension, spouses have nonetheless different treatment status in a large proportion of cases. Furthermore, even when both spouses are treated, the timing of treatment differs for about half of the couples. Panel B shows a very similar picture for wives of treated and non-treated men. Information on exact agreement dates thus allows us to separately identify the direct and cross-effects of shorter workweeks across spouses, as in the majority of cases the year of treatment differs across spouses.

⁸Note that we never use hours reported in the LFS to assign treatment status, but administrative information collected independently by the Ministry of Labour. This prevents us from generating an artificial correlation between our indicator of treatment status and weekly hours.

TABLE 1—DISTRIBUTION OF OWN TREATMENT, BY SPOUSE'S TREATMENT (*percent*)

	Wife not treated	Wife treated
<i>Panel A. Employed men</i>		
Own firm never adopted shorter workweek	71.0	54.2
Own firm adopted shorter workweek	29.0	45.8
Not same year as wife's firm	29.0	22.8
Same year as wife's firm	—	23.0
Total	100	100
	Husband not treated	Husband treated
<i>Panel B. Employed women</i>		
Own firm never adopted shorter workweek	73.2	58.1
Own firm adopted shorter workweek	26.8	41.9
Not same year as wife's firm	26.8	21.3
Same year as wife's firm	—	20.6
Total	100	100

Notes: The sample includes employed respondents. The interpretation of figures is as follows: among employed males whose spouse works in a treated firm, 45.8 percent are working in a treated firm.

Source: French LFS, 1994–2009, INSEE.

B. Graphical Evidence: Direct and Cross-Effect of Treatment

Before moving on to regression analysis, we provide simple graphical evidence on the direct and cross-effects of the workweek reform. Figure 2 plots actual hours worked during the survey week by wives who are wage earners, by treatment status. The solid line refers to treated wives, and time zero refers to the year in which a shorter workweek agreement is implemented at their workplace. Their weekly hours are stable, if anything slightly rising, during the pre-treatment years, and drop by about 2 upon treatment. The dotted line refers to non-treated wives, and reports their working hours for the same dates at which treated wives are observed.⁹ Their weekly hours follow a gradually rising trend throughout the sample period, with no break at time zero. Thus we observe a decline of about two hours in working hours of treated wives relative to control wives at time of treatment. Interestingly, wives who become treated have longer weekly hours initially, and their hours converge almost exactly to hours of non-treated wives when their employers adopt the shorter workweek. Figure B1 (panel A) in the online Appendix plots treatment-control differences in these series, together with the corresponding confidence intervals, and shows flat pre-treatment differences, followed by a permanent, 2-hour drop in correspondence of treatment.

The observed drop in weekly hours for treated wives relative to the non-treated is a first-stage effect for the cross-hour effect on men that we intend to analyze next. A first-stage effect of about 2 hours is equivalent to roughly half the reduction in the legal workweek ($39 - 35 = 4$), and this may be explained by a number of factors,

⁹For each treated individual i , we obtain the average number of hours worked by never treated individuals observed in the LFS in the same year as i , denoted by $H_{c(i)}$. For each $D = -5, -4, \dots, +6$, the dotted line in Figure 2 shows the average of $H_{c(i)}$ over the population of treated individuals i observed at a distance D from treatment (where $D = \text{year of observation} - \text{year of treatment}$).




FIGURE 2. WIVES' HOURS WORKED, BY OWN TREATMENT

including slight redefinitions of working time and/or the possibility to implement the worktime regulation at the annual rather than weekly level¹⁰ (see also Askenazy 2008). This would deliver a mitigated effect of the workweek reduction on mean actual hours in the LFS, as the survey week falls in March of each year, and thus tends not to coincide with popular holiday seasons. Finally, the effect of the introduction of the 35-hour workweek has also been mitigated by the incidence of relatively short workweeks among French employees in the pre-reform period. Specifically, about 39 percent of females and 16 percent of males usually worked less than 39 hours per week before treatment.¹¹ The estimated 2-hour drop in working hours can thus be interpreted as an average of a higher drop for women initially working 39 hours or more, and a smaller drop for those initially working less than 39 hours.

Given the behavior of treated wives, the next question is whether we observe a variation in either the employment rate or the number of hours worked by their husbands. Figure 3 shows flat and virtually identical employment patterns of husbands of treated and non-treated wives. Figure 4 then addresses corresponding variations at the intensive margin, by showing the impact on hours worked by the subsample of employed men, and reveals a sizeable drop in hours worked by husbands of treated women, relative to husbands of non-treated women, at time of treatment. Specifically, the difference in working hours is close to zero during the five pre-treatment years, and rises to 40 minutes on average during the five post-treatment years.

¹⁰ For example, an employer could cut the usual workweek to 37 hours and grant 12.5 additional days of annual leave. In treated firms, about 38 percent of male employees and 23 percent of female employees declare having usual workweeks longer than 35 hours after treatment.

¹¹ Note that for some employees the reform was not even binding, as about 6.5 percent and 31 percent of men and women, respectively, had usual hours below or at 35 in the pre-treatment period. For women, short usual workweeks mostly correspond to part-time work. For men, they correspond mostly to specific jobs and working conditions (e.g., night work, evening work, Sunday work, rotating shift patterns, etc.).




FIGURE 3. MEN'S EMPLOYMENT RATES, BY WIFE'S TREATMENT




FIGURE 4. MEN'S HOURS WORKED, BY WIFE'S TREATMENT

The difference between the two series shows no evidence of differential pre-trends, and jumps permanently upon treatment (online Appendix Figure B2, panel A).

As the observed cross-effects might be partly induced by cases of simultaneous treatment of spouses, we replicate the corresponding trends on a subsample that excludes men treated at the same date as their wives, and on a subsample that excludes men ever treated, respectively. Reassuringly, Figures B3 and B4 in the online Appendix provide a very similar picture of cross-hour effects as Figure 4. In the regression analysis that follows we pool all households and control for own and spouse treatment separately.

Figures 5 to 7 repeat a similar analysis for female respondents and their husbands. Again we observe a clear first-stage effect for husbands (Figure 5), whose




FIGURE 5. HUSBANDS' HOURS WORKED, BY OWN TREATMENT




FIGURE 6. WOMEN'S EMPLOYMENT RATES, BY HUSBAND'S TREATMENT

magnitude is very close to that observed for wives in Figure 2 (differences in these series are plotted in online Appendix Figure B1, panel B). However, we find no evidence of spillover effects on their wives' labor supply, either at the extensive margin (Figure 6), or the intensive margin (Figure 7). The difference between these series is essentially flat, and does not display any permanent jump upon treatment (online Appendix Figure B2, panel B).

The descriptive evidence presented is thus suggestive of labor supply spillovers at the intensive margin for men, but no spillovers at either margin for women. The next sections will show estimates of these effects that control for observable characteristics of the individuals, and explore further the nature of these spillovers.




FIGURE 7. WOMEN'S HOURS WORKED, BY HUSBAND'S TREATMENT

III. Regression Results

A. Main Estimates

As in the previous descriptive analysis, we focus on two main outcome variables for each individual i in our sample, namely her employment status and her weekly hours worked, and assess how each is affected by the implementation of a shorter workweek agreement by her spouse's employer. This would work via an effect on the spouse's labor supply, and thus we start by estimating the first-stage effect of treatment on spouses. We denote by H_{it}^S the actual weekly hours worked by the spouse, and introduce a dummy variable A_{it}^S indicating whether at time t she works for a firm who has ever adopted the shorter workweek. Our first-stage regression is the following difference-in-differences specification:

$$(1) \quad H_{it}^S = \alpha_1 A_{it}^S + \alpha_2 APost_{it}^S + \alpha_3 X_{it}^S + D_t + u_{it},$$

where $APost_{it}^S$ indicates the period following the introduction of the shorter workweek in the spouse's firm, D_t denotes a set of year fixed effects, and X_{it}^S are relevant individual covariates, including a constant term. The α_2 coefficient shows the direct (first-stage) effect of workweek regulations on labor supply.

Table 2 shows the regression results for specification (1) for wives (panel A) and husbands (panel B). All reported standard errors in this and later tables are clustered at the year \times treatment level (32 clusters). Column 1 in panel A shows that wives working in firms who implemented a workweek reduction agreement cut their labor supply by about 1.81 hours per week once the shorter workweek is implemented, as is also evident from Figure 2. Turning to husbands, column 1 in panel B shows again strong and significant effects of the workweek reduction (-1.95 hours). All these estimates are robust to the introduction of controls for age, education, and industry effects (column 2), suggesting that the implementation of the shorter workweek was

TABLE 2—FIRST-STAGE REGRESSIONS:
DIRECT EFFECTS OF THE SHORTER WORKWEEK ON HOURS AND EARNINGS

	Wives' hours		Wives' (log) earnings	
	(1)	(2)	(3)	(4)
<i>Panel A. Men</i>				
<i>A</i> Post ^S	−1.81*** (0.13)	−1.91*** (0.10)	0.002 (0.010)	−0.002 (0.006)
Additional controls	No	Yes	No	Yes
Mean dep. variable	30.05	30.05	8.658	8.658
Observations	189,894	189,894	160,046	160,046
	Husbands' hours		Husbands' (log) earnings	
<i>Panel B. Women</i>				
<i>A</i> Post ^S	−1.95*** (0.13)	−1.92*** (0.14)	0.017** (0.008)	0.007 (0.004)
Additional controls	No	Yes	No	Yes
Mean dep. variable	37.07	37.07	9.011	9.011
Observations	236,802	236,802	201,559	201,559

Notes: The table shows results from first-stage regressions for hours and earnings of spouses. Columns 1 and 2 refer to the full sample (married or cohabiting respondents whose spouse is an employee). Columns 3 and 4 refer to the subsample whose spouses have nonmissing earnings (from 2003 onwards, information on earnings is collected on one-third of the LFS sample). Baseline controls include *A*^S, 15 year dummies, and a dummy for public sector. Additional controls include years of education, age, age squared, and 16 industry dummies. Standard errors clustered at the treatment × year level are reported in brackets.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

Source: French LFS, 1994–2009, INSEE.

largely orthogonal to these job and worker characteristics. Columns 3 and 4 in each panel report estimates of a similar specification for (the log of) monthly earnings, and once extra controls are included these show near zero effects of the workweek reduction on the earnings of wives and husbands. These first-stage results are clearly in line with the reform's intended outcome to shorten the workweek without cutting monthly earnings of treated employees.

We next assess labor supply spillovers by looking at the reduced-form effects of one's spouse's workweek reduction on own employment status and weekly hours. Note that we can interpret such cross-effects as stemming from the sole reduction in the amount of time spent at work by the spouse once we have ruled out the presence of income effects, as shown in columns 3 and 4 of Table 2. Our reduced-form specification for hours is

$$(2) \quad H_{it} = \gamma_1 A_{it}^S + \gamma_2 APost_{it}^S + \gamma_3 A_{it} + \gamma_4 APost_{it} + \gamma_5 X_{it} + D_t + \varepsilon_{it},$$

where H_{it} denotes own weekly hours, A_{it} is a dummy variable denoting whether one's employer has ever implemented a shorter workweek agreement, whereas $APost_{it}$ indicates the period following this agreement. The main coefficient of interest is γ_2 . Note that this specification allows us to estimate cross-effects in labor supply (captured by $APost_{it}^S$), over and above the direct effect of own treatment (captured by $APost_{it}$). These two effects can be separately identified insofar as treatment is not simultaneous for all spouses. A similar linear specification to model (2) is used for

TABLE 3—REDUCED-FORM REGRESSIONS: CROSS-EFFECTS OF THE SHORTER WORKWEEK ON EMPLOYMENT AND HOURS

	Own employment		Own hours (conditional on employment)			
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Men</i>						
$APost^S$	-0.0037 (0.0027)	-0.0028 (0.0022)	-0.44*** (0.09)	-0.45*** (0.09)	-0.50*** (0.09)	-0.44*** (0.10)
$APost$	—	—	-1.96*** (0.14)	-1.96*** (0.14)	-2.02*** (0.13)	—
Further controls	No	Yes	No	Yes	Yes	Yes
Mean dep. variable	0.8819	0.8819	38.89	38.89	38.97	39.55
Observations	189,894	189,894	167,460	167,460	156,392	115,445
	Own employment		Own hours (conditional on employment)			
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel B. Women</i>						
$APost^S$	-0.0032 (0.0023)	-0.0041 (0.0022)	0.12 (0.10)	0.05 (0.11)	0.06 (0.11)	0.07 (0.11)
$APost$			-1.86*** (0.17)	-1.88*** (0.15)	-1.86*** (0.18)	—
Further controls	No	Yes	No	Yes	Yes	Yes
Mean dep. variable	0.6786	0.6786	30.32	30.32	30.25	30.04
Observations	236,802	236,802	160,689	160,689	150,371	116,596

Notes: The table shows results from reduced-form regressions in which own employment status and hours are regressed on spousal treatment A^S (and $APost^S$), as well as on own treatment (A and $APost$). Columns 1 and 2 refer to the full sample. Columns 3 and 4 refer to the subsample of employed respondents. Column 5 refers to employed respondents who were not treated at the same time as their spouses. Column 6 refers to employed respondents who were never treated. Baseline controls in columns 1 and 2 include A^S , 15 year dummies, and spouse's public sector dummy. Additional controls in column 2 are own years of education, age and age squared, and spouse's years of education, age and age squared. Baseline controls in columns 3–6 include A , A^S , 15 year dummies, own public sector and wage-earner dummies, and a spouse's public sector dummy. Additional controls in columns 4–6 include own years of education, age, age squared, and 16 industry dummies, and spouse's years of education, age, age squared, and 16 industry dummies. Standard errors clustered at the treatment \times year level are reported in brackets.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

Source: French LFS, 1994–2009, INSEE.

the extensive margin, where the dependent variable is a dummy for the respondent's employment status, and clearly A_{it} and $APost_{it}$ are not defined.

The regression results are reported in Table 3. Columns 1 and 2 refer to employment, and columns 3–6 refer to weekly hours. Estimates show no evidence of any significant cross-effects on employment for men, and the associated point estimate is always very close to zero, in line with the trends reported in Figure 3. For women, the cross-effect on employment becomes marginally significant when further controls are included in column 2, but its magnitude is negligible. As we find virtually no impact on employment, we next look at hours worked for those who are employed. In column 3 of panel A we regress men's hours on own treatment (A_{it} and $APost_{it}$), and on their wives' treatment (A_{it}^S and $APost_{it}^S$). The own treatment effect is about -2 , and the cross-effect is -0.44 and highly significant, showing that when their wives become subject to the shorter workweek, men reduce their weekly labor supply by nearly half an hour. The magnitude of the cross-effect stays

unchanged when we control for individual characteristics (column 4), and when we exclude men who are treated in the same year as their wives (column 5) or men who are ever treated (column 6). We next let the effect of treatment vary over time, and in particular we estimate a reduced-form specification that includes all controls as in column 4 of Table 3, having interacted A_{it}^S with a full set of pre- and post-treatment dummies. The associated estimates are reported in Figure B5 (panel A) of the online Appendix, and show no pre-treatment effects, together with a permanent drop at time of treatment. In other words, post-treatment estimates are stable and all quite close to the overall treatment effect of -0.44 .

Panel B of Table 3 reports corresponding estimates for women. While the own effect of workweek regulations is negative and significant, the cross-effect is positive, small, and not significantly different from zero. We thus detect no evidence of spousal spillovers in the labor supply of women.

We further explore cross-effects by estimating reduced-form specifications across the whole hours distribution. Specifically, for each k between 15 and 49, we estimate reduced-form equations for the probability of working longer than k hours. These coefficients are reported in Figure 8, together with the corresponding 95 percent confidence interval. For men, cross-effects on hours feature among the whole hours distribution, but most heavily for men working 35–38 hours, and this result replicates very closely on a subsample that excludes men ever treated (graph not reported). For women, cross-effects are much weaker and typically not statistically significant across the entire distribution, but if anything they involve a slight reduction in the incidence of long workweeks ($40 \leq H \leq 45$).

B. Further Estimates: Cross-Effects on Usual and Nonusual Working Hours

We next investigate the nature of labor supply spillovers in further detail by combining information on actual hours (H) with information—also contained in the LFS—on usual hours (H_u), defined as the number of hours worked in a typical week. Actual hours H are the sum of the usual workweek H_u and a nonusual labor supply component $H - H_u$, which may be either positive or negative, depending on whether overtime hours exceed various forms of “undertime” hours (e.g., unusually short working days, sickness absence, paid or unpaid leaves, etc.) in a given week.¹² A worker may reduce weekly hours H by either negotiating a new contract with her employer, involving lower H_u , or keeping her contract unchanged, together with the associated H_u , but cutting on $H - H_u$, and namely some form of work involvement that is typically not specified in a contract. This may imply a reduction in overtime work or an increase in the take-up rate of leaves or in absenteeism. It is reasonable to expect that cross-effects mostly occur through reductions in $H - H_u$, since these would not require the renegotiation of one’s labor contract, and are more easily under an employee’s individual control than adjustments in H_u . On the other hand,

¹² Note that H and H_u represent weekly-aggregated measures, thus someone who works one hour longer than the typical workday for three days in a week and one hour shorter for the remaining two days would have $H > H_u$. For simplicity, we will refer to cases in which $H > H_u$ as cases of overtime work, and to cases in which $H < H_u$ as cases of undertime. Descriptive statistics on overtime and undertime are reported in Section D of the online Appendix.

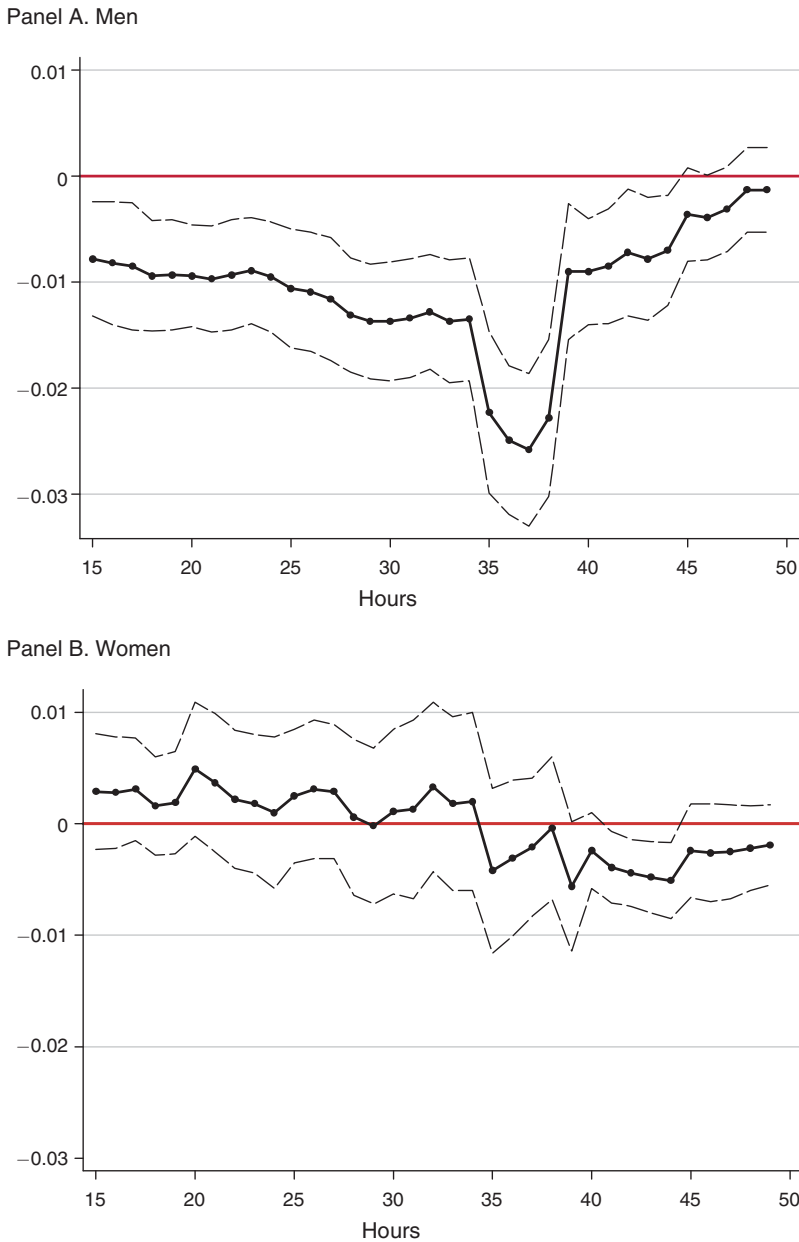


FIGURE 8. ESTIMATED CROSS-EFFECTS ON THE CUMULATIVE DISTRIBUTION OF HOURS

Notes: For each k between 15 and 49, the solid lines show the cross-effect on the probability of working longer than k , i.e., $\Pr(H > k)$. Dashed lines show the corresponding 95 percent confidence intervals.

the direct effect of the law is expected to bite on H_u , consistently with the collective nature of these agreements.

Estimates reported in Table 4 shed light on these adjustment margins. The sample period is now restricted to 1994–2002, as information on usual hours is unavailable from 2003 onwards. Estimates in panel A refer to men. Columns 1 and 2 show that, as anticipated, the first-stage effect of the workweek reduction in their wives’ firms

TABLE 4—FIRST-STAGE AND REDUCED-FORM REGRESSIONS: DIRECT AND CROSS-EFFECTS OF THE SHORTER WORKWEEK ON USUAL AND NONUSUAL HOURS

	First-stage		Reduced-form	
	Wife's usual hours H_u (1)	Wife's actual-usual hours $H - H_u$ (2)	Own usual hours H_u (3)	Own actual-usual hours $H - H_u$ (4)
<i>Panel A. Men</i>				
A_{Post}^S	-1.75*** (0.15)	-0.54*** (0.16)	-0.05 (0.05)	-0.62*** (0.14)
Mean dep. var.	33.79	-4.46	39.24	-3.17
Observations	97,470	97,470	97,470	97,470
	First-stage		Reduced-form	
	Husband's usual hours H_u (1)	Husband's actual-usual hours $H - H_u$ (2)	Own usual hours H_u (3)	Own actual-usual hours $H - H_u$ (4)
<i>Panel B. Women</i>				
A_{Post}^S	-2.02*** (0.12)	-0.46 (0.23)	0.17** (0.08)	0.06 (0.10)
Mean dep. var.	39.17	-3.18	33.33	-4.28
Observations	102,123	102,123	102,123	102,123

Notes: Regressions refer to the employed subsample with nonmissing own and spouse's usual hours. Control variables in columns 1 and 2 are the same as in column 2 of Table 2, and in columns 3 and 4 they are the same as in column 4 of Table 3. Standard errors clustered at the treatment \times year level are reported in brackets.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

Source: French LFS, 1994–2002, INSEE.

mostly bites on usual hours (-1.75), while the effect on nonusual hours is much weaker (-0.54). By contrast, columns 3 and 4 show that the reduced-form effect of the reform on own hours works entirely via a reduction in nonusual hours (-0.62), with no cross-effect on usual hours (-0.05), and thus no need to renegotiate own work schedules for men responding to their wives' work schedules. For women (panel B), we detect very similar first-stage effects as for men, but a small, albeit positive, cross-effect on H_u (0.17).

Changes in nonusual hours and earnings are further explored in Table 5. Columns 1 and 2 report cross-hour effects on overtime hours and undertime hours separately. These are defined as $(H - H_u)^+ = \max(H - H_u, 0)$ and $(H - H_u)^- = \max(H_u - H, 0)$, respectively. Cross-hour effects feature strongly on undertime hours (0.54), while overtime hours are hardly affected (-0.07). Cross-effects on undertime hours in turn involve an increase in the frequency of both unworked weeks ($H = 0$, column 3) and unusually short workweeks ($0 < H < H_u$, column 4), but no change at all in full-time status (column 5). For cases in which $H < H_u$, respondents are asked whether they worked less than usual in the reference week due to holidays and absence for personal reasons, sickness leave, maternity leave, continuous training, unusual workload, strike, or lock-out. While we detected significant cross-hour effects for holidays and sickness leaves, which are margins on

TABLE 5—REDUCED-FORM REGRESSIONS: CROSS-EFFECTS OF THE SHORTER WORKWEEK ON TYPES OF HOURS WORKED AND EARNINGS

	Own overtime hours ($H - H_U$) ⁺ (1)	Own undertime hours ($H - H_U$) ⁻ (2)	Own unworked weeks $H = 0$ (3)	Own unusually short workweeks $0 < H < H_U$ (4)	Own part-time (5)	Own (log) earnings (6)
<i>Panel A. Men</i>						
<i>A</i> Post ^s	-0.07 (0.03)	-0.54*** (0.11)	0.012*** (0.003)	0.006** (0.003)	-0.002 (0.002)	0.002 (0.003)
Mean dep. var.	0.86	-4.03	0.088	0.065	0.031	9.004
Observations	97,470	97,470	97,470	97,470	97,470	97,470
<i>Panel B. Women</i>						
<i>A</i> Post ^s	-0.06 (0.03)	0.12 (0.10)	-0.005 (0.003)	0.005 (0.003)	-0.012*** (0.004)	0.020*** (0.004)
Mean dep. var.	0.56	-4.84	0.129	0.061	0.323	8.587
Observations	102,123	102,123	102,123	102,123	102,123	102,123

Notes: Regressions refer to the employed subsample with nonmissing own and spouse's usual hours. Control variables are the same as in column 4 of Table 3. In column 2, the interpretation of positive coefficients is that the fall in labor supply is now picked up by an increase in undertime hours. Standard errors clustered at the treatment \times year level are reported in brackets.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

Source: French LFS, 1994–2002, INSEE.

which employees have closer control, we found no evidence of cross-effects on any other margin (results not reported).¹³

Finally, we do not find any detrimental cross-effect on male earnings (column 6), consistently with evidence on the contribution of various components of actual hours (usual, overtime and undertime, respectively) to monthly earnings, as illustrated in Table C1 in the online Appendix. Interestingly, undertime hours turn out to be the sole component of labor supply that men may cut unilaterally without earning losses.

No hours margin is significantly affected for women (panel B), except the incidence of part-time work, which falls by nearly 1 percentage point. The slight increase in the usual workweek and the corresponding change in full-time status are accompanied by an increase in earnings (2 percent), in line with the fact that usual hours are the labor supply component that best predicts earnings (online Appendix Table C1).

In summary, we detect substantial differences in both the magnitude and nature of spillover effects across genders. Specifically, cross-effects do not entail the renegotiation of usual hours with employers or changes in earnings for men, but involve instead a reduction in their unusual work involvement, whether within a given day, or through an increase in the take-up rate of paid vacation or sick leave, with no detrimental impact on (current) earnings. A reason why men may work some unpaid hours in the first place is that these may have an impact on future, as opposed to current, earnings, to the extent that someone who is more absent from work may

¹³ Information on the take-up rate of paid leaves and paid and unpaid overtime work contained in later waves of the LFS (2003–2009) confirms that there exists significant leeway for most employees, and especially for the high-skilled, in reducing their unpaid involvement at work.

lose on prospects of promotion and/or earnings growth. Another possible explanation is that some individuals may derive utility from work per se. Regardless of the underlying mechanism, our results show that men decide to cut on such unpaid hours following their wives' treatment, as increased spousal non-market time would raise the utility of their own non-market time relative to the utility of being at work.

Women, by contrast, are more often working part-time and less often spending unpaid, nonusual hours at work. Compared to men, it is on average more costly for women to adjust hours downward, insofar as they have lower nonusual hours margins than men, but less costly to adjust hours upward, as in the public sector and large private sector firms employees can easily shift from part-time to full-time status, and only among women is the incidence of part-time work substantial. The French reform thus provides a clean example of the role of optimization frictions in shaping the magnitude and nature of social spillovers.

C. Heterogeneous Cross-Hour Effects

As working hours, constraints, and preferences may vary widely across individuals, cross-hour effects may differ across occupations and the household composition of workers. Workers in high-skill occupations (managers, professionals, and associate occupations) on average work longer hours than the less-skilled and typically have higher control over the organization of their workweek, while the less-skilled are more likely to work the legal workweek and thus would only be able to cut their working hours via new contractual agreements.

Panels A and B in Table 6 replicate our previous analysis on actual hours for employees in high-skill occupations and other employees, respectively. First-stage effects reported in column 1 have conventional magnitude and significance. For men, the associated cross-effect on hours is about three times larger for high-skill occupations (column 2) than for other occupations (column 3). Similar conclusions can be drawn by looking at the probability of working more than 45 hours weekly (columns 3 and 6). Spillover effects on men's labor supply thus seem much stronger for the high-skilled than for the less-skilled.¹⁴ For women, we do not find significant cross-effects on overall working hours, but we do find a negative and significant impact on the probability that females in high-skill occupations work long weeks. This is the only subsample and only outcome variable for which we detect symmetric cross-effects for men and women. We found in Section IIIA that women are slightly less likely to work very long hours when their husbands are treated (Figure 8), and we note here that for female managers and professionals this effect is as strong as for men, suggesting that when women have enough leeway to cut their hours—either because they work very long hours in the first place or they have managerial control—their labor supply response is qualitatively similar to that of men. However, the subsample of such women is too small, and their labor supply response too weak, for this effect to be discernible on the full sample.

¹⁴ In the online Appendix, we also show that cross-effects for men are stronger in the public than in the private sector, consistently with the presumption that public employees in France tend to have, other things equal, greater control than private employees in organizing their working time (see Table D1, panel A).

TABLE 6—HETEROGENEOUS EFFECTS OF THE SHORTER WORKWEEK, BY OCCUPATION AND FAMILY TYPE

	Managers, profs., and kindred occup.			Other occupations		
	First-stage	Reduced-form		First-stage	Reduced-form	
	Wife's hours (1)	Own hours (2)	Own hours ≥ 45 (3)	Wife's hours (4)	Own hours (5)	Own hours ≥ 45 (6)
<i>Panel A. Men</i>						
<i>A</i> Post ^S	-2.32*** (0.30)	-0.81*** (0.27)	-0.033*** (0.009)	-1.72*** (0.11)	-0.32*** (0.09)	-0.006*** (0.002)
Mean dep. var	29.80	40.91	0.447	30.20	38.44	0.217
Observations	30,432	30,432	30,432	137,028	137,028	137,028
	Managers, profs., and kindred occup.			Other occupations		
	First-stage	Reduced-form		First-stage	Reduced-form	
	Husband's hours (1)	Own hours (2)	Own hours ≥ 45 (3)	Husband's hours (4)	Own hours (5)	Own hours ≥ 45 (6)
<i>Panel B. Women</i>						
<i>A</i> Post ^S	-2.51*** (0.40)	-0.17 (0.34)	-0.034*** (0.008)	-2.03*** (0.13)	0.15 (0.11)	-0.001 (0.002)
Mean dep. var	38.51	32.03	0.196	36.91	30.14	0.069
Observations	15,217	15,217	15,217	145,472	145,472	145,472
	At least one child aged 0–6			No children aged 0–6		
	First-stage	Reduced-form		First-stage	Reduced-form	
	Wife's hours (1)	Own hours (2)	Own hours ≥ 45 (3)	Wife's hours (4)	Own hours (5)	Own hours ≥ 45 (6)
<i>Panel C. Men</i>						
<i>A</i> Post ^S	-1.30*** (0.23)	-0.81*** (0.28)	-0.028*** (0.008)	-2.08*** (0.13)	-0.34*** (0.12)	-0.003 (0.003)
Mean dep. var	27.53	38.85	0.260	30.93	38.91	0.259
Observations	39,468	39,468	39,468	127,992	127,992	127,992
	At least one child aged 0–6			No children aged 0–6		
	First-stage	Reduced-form		First-stage	Reduced-form	
	Husband's hours (1)	Own hours (2)	Own hours ≥ 45 (3)	Husband's hours (4)	Own hours (5)	Own hours ≥ 45 (6)
<i>Panel D. Women</i>						
<i>A</i> Post ^S	-2.25*** (0.24)	0.49 (0.25)	-0.001 (0.003)	-2.04*** (0.11)	-0.08 (0.10)	-0.005** (0.002)
Mean dep. var	37.16	27.94	0.063	37.03	31.03	0.086
Observations	36,959	36,959	36,959	123,730	123,730	123,730

Notes: Regressions refer to the employed subsample. In columns 1 and 4, control variables are the same as in column 4 of Table 3, and in columns 2, 3, 5, 6 they are the same as in column 4 of Table 4. Standard errors clustered at the treatment × year level are reported in brackets.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

Source: French LFS, 1994–2009, INSEE.

We further explore spousal labor supply responses across household types. It has been argued that interdependences in spousal labor supply may be stronger in the presence of young children, as children appear to play the role of a jointly-consumed commodity for husbands and wives (Lundberg 1988). Panels C and D of Table 6 cover households with at least one child aged 0–6, and other households separately. We find weaker first-stage effects for mothers of young children than for other women, in line with higher incidence of part-time work among mothers, as for part-timers the mandatory workweek reduction is not necessarily binding. Reduced-form regressions show a much stronger labor supply reaction for fathers of young kids than other men, despite a weaker first stage. For women, cross-effects are somewhat mixed, as we detect a positive, rather than negative, cross-hour effect for mothers of young kids, and a negative cross-effect on the probability to work long weeks for other women.

IV. Robustness Analysis

The identifying assumption underlying our main estimates is that a respondent's unobserved characteristics be uncorrelated to the timing of adoption of the shorter workweek in his or her spouse's firm. One may think of scenarios in which this assumption is potentially violated, and we perform a number of robustness tests that should address various caveats to a causal interpretation of our estimates.

First, one should worry about the existence of differential pre-existing trends in working hours between treatment and control groups. However, the event-study type of evidence presented in Figures 2–7 clearly shows that this is not the case, as pre-trends are in all cases parallel or even flat. This is also confirmed by estimates of reduced-form specifications that control for treatment-specific trends, reported in columns 1 and 4 of Table E1 in the online Appendix. Columns 3 and 6 in Table E1 further control for region \times year interactions, capturing the effect of local shocks, and show virtually unchanged estimates from columns 1 and 4, respectively.

Second, our identifying assumption would be violated if spouses of employees in firms adopting the shorter workweek were subject to systematically different shocks or changes in unobservables around time of adoption, versus spouses of employees in non-adopting firms. If changes in unobservables of treatment and control groups would generate spurious changes to their labor supply, one would possibly expect to observe some change in some of their observables as well around the time of treatment. But online Appendix Table E2 shows no evidence of any significant change in such characteristics upon treatment. Third, we take into account concerns of reverse causality, namely the possibility that changes in own labor supply behavior may affect spousal job mobility between adopting and non-adopting firms, and replicate our reduced-form specifications on a subsample of spouses of job-stayers (online Appendix, page 5).

Finally, one may worry that in general employees in adopting (or early-adopting) firms would have systematically different spouses from employees in non-adopting (or late-adopting) firms. To address this concern, we provide fixed-effect estimates of the effects of interest, based on a (limited) rotating panel component of the LFS (online Appendix Table E4). This last robustness test confirms our main estimates.¹⁵

¹⁵ We also checked that our estimates are very similar whether identification only relies on variation in hours across treated and non-treated spouses, or across early and late-treated spouses (Section E.2 in the online Appendix).

V. Instrumental Variable Estimates of Cross-Hour Effects

There is a long standing tradition of labor supply models in which the decisions of each spouse depend on the number of hours spent at work by the other spouse (see Lundberg 1988, for a seminal example). These models are hard to estimate as they involve a system of two simultaneous equations in which wives' hours feature in the husbands' labor supply equation and vice versa, and good instruments for independent variation in the labor supply of one of the spouses are typically hard to find. In such a scenario the French workweek reform helps identify the effects of interest by generating exogenous variation in the labor supply of one's spouse.

While the previous sections have highlighted the reduced-form effect of workweek regulations on spousal labor supply, in this section we use workweek regulations in an individual's firm as an instrument for her working hours in her spouse's labor supply equation. Under the exclusion restriction that workweek regulations affect spousal labor supply only via their effect on the labor supply of directly treated employees, IV estimates provide the parameter of interest for measuring how labor supply responds to independent changes in labor supply of one's spouse, and may be generalized to a variety of scenarios.

The structural interpretation of this parameter, as well as of its variation across genders, relies on the underlying model of intra-household interactions, and in particular on whether one assumes the household decision making process to be cooperative or non-cooperative. In non-cooperative models (see for instance, Bourguignon 1984; Chen and Woolley 2001; Lechene and Preston 2011), each spouse maximizes an individual utility function, taking the decisions of the other spouse as given. The arguments of such utility functions may include own as well as spousal use of time. In this framework cross-hour effects represent the effect of spousal labor supply on the marginal utility of substituting time spent at work with leisure. Asymmetric cross-hour effects can be easily generated in this context by different utility functions for men and women, such that men's utility of leisure would respond to wives' leisure, but not vice versa. In cooperative household models (see, among others, McElroy and Horney 1981; Chiappori 1988; Apps and Rees 1988), the household jointly maximizes a utility function, strictly increasing in the utility of each spouse. In this case it can be shown that estimated cross-hour effects for men and women stem from the same set of parameters in spouses' utility functions and, consequently, strongly asymmetric cross-hour effects for men and women are less straightforward to rationalize, unless women are initially trapped at a corner solution characterized by zero unpaid time at work (see detailed discussion in Section F of the online Appendix).

Below we report estimates of the impact of spousal hours on own hours, having instrumented spousal hours by $APost^S$. The regression results are reported in Table 7 for men (panel A) and women (panel B), using the same samples and specifications as in Tables 3 and 6. Among men, the average cross-hour effect in labor supply is 0.23, but about twice as large for managers and professionals than for other occupations. When their wives cut their labor supply by one hour, men in high occupations respond by cutting their own labor supply by about 20 minutes. Also, cross-effects are three times larger in the presence of young children, relative to other households. The quantitative response for fathers is about 35 minutes for each extra hour spent at home by their wives, suggesting that worktime policy evaluations restricted to

TABLE 7—IV ESTIMATES OF CROSS-HOUR EFFECTS

	Own hours				
	All (1)	High-skilled (2)	Other occupation (3)	One or more child 0–6 (4)	Other households (5)
<i>Panel A. Employed men</i>					
Wife's hours	0.23*** (0.05)	0.34*** (0.12)	0.18*** (0.05)	0.59*** (0.21)	0.16*** (0.06)
Mean dep. variable	38.89	40.91	38.44	38.85	38.91
Observations	167,460	30,432	137,028	39,468	127,992
<i>Panel B. Employed women</i>					
Husband's hours	-0.02 (0.05)	0.08 (0.13)	-0.07 (0.06)	-0.23 (0.12)	0.04 (0.05)
Mean dep. variable	30.32	32.03	30.14	27.94	31.03
Observations	160,689	15,217	145,472	36,959	123,730

Notes: Regressions refers to the employed subsample. Estimates reported show the effect of spousal hours (H^S) on own hours (H), using spousal treatment ($A\text{Post}^S$) as an instrument. The corresponding reduced-form results are reported in Tables 3 and 6. Further controls include A , $A\text{Post}$, A^S , 15 year dummies, a wage-earner dummy and the following variables for each spouse: a public sector dummy, years of education, age, age squared, and 16 industry dummies. Standard errors clustered at the treatment \times year level are reported in brackets.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

Source: French LFS, 1994–2009, INSEE.

direct labor supply effects may strongly underestimate its impact on the time spent by fathers with their young children. For women we detect no significant cross-hour effect on the whole sample or across the occupational divide, but we do find a negative, marginally significant cross-effect for mothers of young kids.

These estimates can be used to quantitatively evaluate the social multiplier, i.e., the gap between aggregate and individual effects of a labor supply shock. Macroeconomic calibrations existing in the literature typically imply much higher labor supply elasticities than individual-level estimates (Chetty et al. 2011a, b), and spousal labor supply complementarities represent an important channel for such a gap. Our estimates reveal a strongly asymmetric structure of spillovers, whereby women's treatment affects male labor supply but not vice versa (with very few exceptions). Specifically, an average cross-hour effect of 0.23 for husbands and a negligible one for wives means that a unit change in individual hours translates into a change in household labor supply of 2.23. This implies a macro response that is $2.23/2 - 1 = 11.5$ percent higher than the micro response for the average household. As discussed by Glaeser, Sacerdote, and Scheinkman (2003), the role of social interactions and social multipliers may vary widely across demographic groups and levels of aggregation, and the French workweek reform provides a clean experiment to identify the multiplier in labor supply at the household level.

Finally, our findings on specific margins of adjustments of weekly hours reveal that, due to search frictions and hours constraints, it is mostly nonusual hours that respond to spouse treatment, leaving usual hours largely unchanged. Thus the above estimates of the social multiplier are likely attenuated by optimization frictions, and may be interpreted as a lower bound for macro elasticities that one would observe absent frictions (Chetty et al. 2011a, Chetty 2012).

VI. Conclusions

We have investigated cross-hour effects in the labor supply of couples using independent variation in spousal hours generated by changes in worktime regulations. In particular we exploit independent variation in spousal hours at constant monthly earnings, which allows us to abstract from income effects of changes in spousal labor supply, and focus on pure cross-hour effects. While wives of treated men hardly adjust their working time, husbands of treated women respond by cutting their workweek by about half an hour to one hour, according to specifications and samples. Such gender differences in cross-hour effects are remarkable; especially insofar as women's labor supply elasticity is typically higher than men's (Blundell and MaCurdy 1999). These results suggest significant spousal complementarities in leisure time for men. While we do not find strong evidence on different preferences by gender, insofar as women work shorter hours in the first place and are less likely than men to have managerial control, they may be more heavily constrained in the organization of their working time.

Our results on cross-hour effects are noteworthy as they show that neglecting spousal responses may give a misleading view of the overall impact of labor supply shocks. In particular, evaluations restricted to the direct impact of policy on the targeted population are likely to underestimate its overall effect on labor supply. A simple back-of-envelope calculation suggests a social multiplier around 1.11, thus neglecting spillovers within the household would yield an underestimate of the overall policy impact on labor supply by about 11 percent. Finally, cross-hour effects vary widely across household types, and tend to be strongest in the presence of young children, with policy relevant effects on the time spent by fathers with their offspring.

REFERENCES

- Abowd, John M., and Francis Kramarz.** 1999. "The Analysis of Labor Markets Using Matched Employer-Employee Data." In *Handbook of Labor Economics*, Vol. 3B, edited by Orley C. Ashenfelter and David Card, 2629–2710. Amsterdam: Elsevier.
- Alesina, Alberto, Edward Glaeser, and Bruce Sacerdote.** 2006. "Work and Leisure in the United States and Europe: Why so Different?" In *NBER Macroeconomics Annual 2005*, Vol. 20, edited by Mark Gertler and Kenneth Rogoff, 1–100. Cambridge, MA: MIT Press.
- Apps, Patricia F., and Ray Rees.** 1988. "Taxation and the Household." *Journal of Public Economics* 35 (3): 355–69.
- Ashenfelter, Orley, and James J. Heckman.** 1974. "The Estimation of Income and Substitution Effects in a Model of Family Labor Supply." *Econometrica* 42 (1): 73–85.
- Askenazy, Philippe.** 2008. "A Primer on the 35-Hour in France, 1997–2007." Institute for the Study of Labor Discussion Paper 3402.
- Baum-Snow, Nathaniel, and Derek Neal.** 2009. "Mismeasurement of Usual Hours Worked in the Census and ACS." *Economics Letters* 102 (1): 39–41.
- Blau, David M.** 1998. "Labor Force Dynamics of Older Married Couples." *Journal of Labor Economics* 16 (3): 595–629.
- Blundell, Richard, and Thomas MaCurdy.** 1999. "Labor Supply: A Review of Alternative Approaches." In *Handbook of Labor Economics*, Vol. 3A, edited by Orley Ashenfelter and David E. Card, 1559–1695. Amsterdam: Elsevier Science.
- Bourguignon, François J.** 1984. "Rationalité individuelle ou rationalité stratégique: le cas de l'offre familiale de travail." *Revue Économique* 35 (1): 147–62.
- Chemin, Matthieu, and Etienne Wasmer.** 2009. "Using Alsace-Moselle Local Laws to Build a Difference-in-Differences Estimation Strategy of the Employment Effects of the Thirty-Five-Hour Workweek Regulation in France." *Journal of Labor Economics* 27 (4): 487–524.
- Chen, Zhiqi, and Frances Woolley.** 2001. "A Cournot-Nash Model of Family Decision Making." *Economic Journal* 111 (474): 722–48.

- Chetty, Raj.** 2012. "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply." *Econometrica* 80 (3): 969–1018.
- Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri.** 2011a. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *Quarterly Journal of Economics* 126 (2): 749–804.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber.** 2011b. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review* 101 (3): 471–75.
- Chiappori, Pierre-André.** 1988. "Rational Household Labor Supply." *Econometrica* 56 (1): 63–90.
- Crépon, Bruno, and Francis Kramarz.** 2002. "Employed 40 Hours or Not Employed 39: Lessons from the 1982 Mandatory Reduction of the Workweek." *Journal of Political Economy* 110 (6): 1355–89.
- Cullen, Julie Berry, and Jonathan Gruber.** 2000. "Does Unemployment Insurance Crowd Out Spousal Labor Supply?" *Journal of Labor Economics* 18 (3): 546–72.
- Estevo, Marcello, and Filipa Sa.** 2008. "The 35-Hour Workweek in France: Straightjacket or Welfare Improvement?" *Economic Policy* 23 (55): 417–59.
- Gelber, Alexander M.** Forthcoming. "Taxation and the Earnings of Husbands and Wives: Evidence from Sweden." *Review of Economics and Statistics*.
- Glaeser, Edward L., Bruce I. Sacerdote, and Jose A. Scheinkman.** 2003. "The Social Multiplier." *Journal of the European Economic Association* 1 (2–3): 345–53.
- Goux, Dominique, and Eric Maurin.** 1999. "Persistence of Interindustry Wage Differentials: A Reexamination Using Matched Worker-Firm Panel Data." *Journal of Labor Economics* 17 (3): 492–533.
- Goux, Dominique, Eric Maurin, and Barbara Petrongolo.** 2014. "Worktime Regulations and Spousal Labor Supply: Dataset." *American Economic Review*. <http://dx.doi.org/10.1257/aer.104.1.252>.
- Gustman, Alan L., and Thomas L. Steinmeier.** 2000. "Retirement in Dual-Career Families: A Structural Model." *Journal of Labor Economics* 18 (3): 503–45.
- Hamermesh, Daniel S.** 2002. "Timing, Togetherness, and Time Windfalls." *Journal of Population Economics* 15 (4): 601–23.
- Hunt, Jennifer.** 1998. "Hours Reductions as Work-Sharing." *Brookings Papers on Economic Activity* (1): 339–69.
- International Labour Office.** 2002. *Key Indicators of the Labour Market: 2001–2002*. New York: Routledge.
- Lechene, Valérie, and Ian Preston.** 2011. "Noncooperative Household Demand." *Journal of Economic Theory* 146 (2): 504–27.
- Lee, Sangheon, Deirdre McCann, and Jon C. Messenger.** 2007. *Working Time around the World: Trends in Working Hours, Laws, and Policies in a Global Comparative Perspective*. New York: Routledge.
- Lundberg, Shelly.** 1985. "The Added Worker Effect." *Journal of Labor Economics* 3 (1): 11–37.
- Lundberg, Shelly J.** 1988. "Labor Supply of Husbands and Wives: A Simultaneous Equations Approach." *Review of Economics and Statistics* 70 (2): 224–35.
- Maurin, Eric, and Julie Moschion.** 2009. "The Social Multiplier and Labor Market Participation of Mothers." *American Economic Journal: Applied Economics* 1 (1): 251–72.
- McElroy, Marjorie B., and Mary Jean Horney.** 1981. "Nash-Bargained Household Decisions: Toward a Generalization of the Theory of Demand." *International Economic Review* 22 (2): 333–49.