

IZA DP No. 2207

Biological Gender Differences, Absenteeism and the Earning Gap

Andrea Ichino
Enrico Moretti

July 2006

Biological Gender Differences, Absenteeism and the Earning Gap

Andrea Ichino

*University of Bologna, EUI,
CEPR, CESifo and IZA Bonn*

Enrico Moretti

*University of California, Berkeley
and IZA Bonn*

Discussion Paper No. 2207
July 2006

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
Email: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of the institute. Research disseminated by IZA may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit company supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its research networks, research support, and visitors and doctoral programs. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Biological Gender Differences, Absenteeism and the Earning Gap^{*}

In most Western countries illness-related absenteeism is higher among female workers than among male workers. Using the personnel dataset of a large Italian bank, we show that the probability of an absence due to illness increases for females, relative to males, approximately 28 days after a previous illness. This difference disappears for workers age 45 or older. We interpret this as evidence that the menstrual cycle raises female absenteeism. Absences with a 28-day cycle explain a significant fraction of the male-female absenteeism gap. To investigate the effect of absenteeism on earnings, we use a simple signaling model in which employers cannot directly observe workers' productivity, and therefore use observable characteristics – including absenteeism – to set wages. Since men are absent from work because of health and shirking reasons, while women face an additional exogenous source of health shocks due to menstruation, the signal extraction based on absenteeism is more informative about shirking for males than for females. Consistent with the predictions of the model, we find that the relationship between earnings and absenteeism is more negative for males than for females. Furthermore, this difference declines with seniority, as employers learn more about their workers' true productivity. Finally, we calculate the earnings cost for women associated with menstruation. We find that higher absenteeism induced by the 28-day cycle explains 11.8 percent of the earnings gender differential.

JEL Classification: J7, M5

Keywords: absenteeism, gender differentials

Corresponding author:

Andrea Ichino
European University Institute
Via della Piazzola 43
Firenze, 50133
Italy
E-mail: andrea.ichino@iue.it

^{*} We would like to thank Sascha Becker, Rebecca Blank, Daniele Checchi, Claudia Goldin, Stefano Gagliarducci, Bryan Graham, Soren Johansen, David Lee, David Levine, Lisa Lynch, Eliana La Ferrara, Massimo Motta, Michele Pellizzari, Steve Pischke, Tuomas Pekkarinen, Simonetta Salvini, Betsey Stevenson, Marko Trevio, Daniela Vuri, Walter Willet and seminar participants at Berkeley, Bologna, CESifo, EUI, Houston, IFAU, IHS Vienna, Linz, Milano, NBER Personnel Meeting, NBER Labor Meetings, Paris, Public Policy Institute of California, Rice, Siena, SOLE Meetings, Texas A&M and Torino for insightful discussions and useful suggestions. We are grateful to Christine Neill for estimating one of the regressions in Table 1 on restricted Canadian data. Jane Leber Herr and Ashley Langer provided excellent research assistance.

1 Introduction

In most Western countries absenteeism is higher among female workers than among male workers. Column 1 of Table 1 shows that in Europe, women take approximately 6.7 more sick-days per year than men. This number includes only illness-related absences, and therefore excludes maternity leave. In the US and Canada, the corresponding figures are 3 and 5.2 days. If we control for age, education and occupation, these differences do not decline (column 2). Furthermore, family-related commitments can explain only part of this gender gap in illness-related absenteeism. For instance, when we restrict the comparison to unmarried workers with no children, we see that in Europe women still take almost 3 more sick-days than men (column 4). The corresponding figures for the US and Canada are 2 and 1.1 days.¹

In this paper we argue that part of this gender difference in absenteeism may be attributed to a biological difference between men and women and that this biological difference has non-trivial earnings consequences for women.

Using the personnel dataset of a large Italian bank, which contains the exact date and duration of every employee absence from work, we find that the hazard of an absence due to illness increases significantly for females, relative to males, 28 days after the previous absence. While the gender difference in hazard is large for those 45 or younger, there is no evidence of such a difference for older employees. We interpret this evidence as suggesting that the menstrual cycle increases women's absenteeism. Absences with 28-day cycles are an important determinant of gender differences in sick days, explaining roughly a third of the overall gender gap in days of absence, and more than two-thirds of the overall gender gap in the number of absences. Our estimate of the incidence of menstrual symptoms are consistent with the existing medical literature.

Furthermore, the incidence of the observed 28-day cycle does not seem to be systematically correlated with employee incentives or local social norms. For example, we find that the cycle is no less pronounced for managers than for clerks, even though the former are less likely to shirk and have significantly fewer absences overall. Similarly, the cycle is no less pronounced for those workers up for promotion, who arguably have strong incentives to

¹We are not the first to document that women have higher levels of absenteeism than men. See, for example, Paringer (1983), Leigh (1983), Barmby et. al (1991), VandelHeuvel and Wooden (1995), Vistnes (1997), and Bridges and Mumford (2000). The literature has not provided convincing evidence on what the causes and consequences of these gender differences may be.

minimize shirking. In fact, the cycle is slightly more pronounced in the months leading up to a promotion than in the months immediately following, even though overall absenteeism rises after a promotion. We also find that, although there is enormous variation in the level of total absenteeism across bank branches, the incidence of the 28-day cycle is not correlated with the local level of total absenteeism.

What is the effect on women’s earnings and careers of this additional absenteeism? In the second part of this paper we present a simple model that clarifies how the relationship between absenteeism, earnings and worker quality may differ for men and women. In particular, we argue that an important component of the effect of an absence on earnings arises from its signaling value. In the model, employers cannot directly observe individual productivity. Instead they use observable worker characteristics—including absenteeism—to predict productivity and set wages.

The key insight of the model is that absenteeism is a noisier signal of worker quality for females than for males. While male absenteeism depends only on the propensity to shirk and on non-menstrual health shocks, female absenteeism is also driven by the menstrual cycle. Signal extraction of underlying shirking rates based on absenteeism is therefore more informative for men than for women. As a result, the relationship between earnings and absenteeism should be more negative for men. A second implication is that this gender difference in the slope between earnings and absenteeism should decline with seniority. As employers learn more about a worker’s true productivity, the importance of the signal should decline.²

Our data seem remarkably consistent with the predictions of this model. First, we find that the relationship between earnings and cyclical absenteeism is negative for both genders, with the slope significantly steeper for men. In other words, the cost of a day’s absence is lower for women. Second, we find the same difference in slope when we look at the relationship between absenteeism and other indicators of worker quality, such as education or the number of episodes of misconduct. Third, this gender difference in slope is large when an employee first joins the firm, and declines with seniority. Consistent with the notion that employers learn about workers’ productivity over time, the negative relationship between earnings and absenteeism is the same for those men and women with 15 years’ seniority.

²These predictions remain true in a model where workers can endogenously choose their effort level to reduce absenteeism.

In addition to showing higher overall rates of absenteeism, women in this sample earn about 13.5% less than men, conditional on their demographic characteristics. This difference is similar to what we observe in representative samples of white collar workers from the US or Europe. In the final part of this paper we calculate how much of this gender gap in earnings can be attributed to the additional absenteeism induced by the menstrual cycle. To do this, we construct a counterfactual earnings gap in the absence of menstruation by assigning the male distribution of absenteeism to females, and re-weighting the conditional earnings gap based on these counterfactual weights.

We find that in the absence of 28-day cyclical absenteeism, the conditional gender gap in earnings would decline from -13.5% to -11.9%, an 11.8 percent decline. About a third of this effect is explained by the direct loss of output associated with additional absenteeism induced by the menstrual cycle. The remaining two-thirds are explained by signaling and other costs. Absenteeism associated with the 28-day cycle explains an even larger fraction of the gender gap in careers. In particular, it explains 13.6 percent of the gender gap in the probability of promotion to management. These counterfactual calculations should be interpreted as lower bounds of the effect of menstrual episodes, since according to our model, the decline in worker quality associated with increases in absenteeism should be more pronounced for men than for women.

Our findings may have policy implications. Forcing employers, rather than women, to bear the monetary burden associated with menstruation may be counterproductive. Whether society should address this biological difference with a gender-based wage subsidy depends on voters' tastes for redistribution. Clearly, this is not a case of market failure, and the rationale for the subsidy would be redistribution rather than efficiency. A gender-specific public subsidy financed out of general taxation would shift part of the costs of menstrual-related absenteeism from women to men.

The paper proceeds as follows. In Section 2 we test whether menstrual symptoms increase women's absenteeism. In Section 3 we then use the predictions of a simple model of wage determination to investigate how the cost of an absence varies between men and women. In Section 4 we quantify how much of the gender gap in earnings can be explained by the additional absenteeism induced by the menstrual cycle, and in Section 5 we conclude.

2 Is There a 28-Day Cycle in Female Absenteeism?

In this Section we test whether women’s absences from work display a systematic 28-day cycle. We begin by showing some graphical evidence (subsection 2.1) and then present more formal parametric estimates (subsection 2.2). We find that work absences of younger women do display an approximate 28-day cycle, and that this cycle disappears for women over age 45. We then quantify the number of absences associated with this cycle. We find that the 28-day pattern is responsible for a significant fraction of the gender gap in absenteeism (subsection 2.3). Finally, we show that this 28 day cycle is apparent across different types of women. In particular, it affects managers as well as clerks, and women who are up for promotion, as well as those who are not (subsection 2.4).

2.1 Graphical Evidence

We use a dataset comprised of personnel data for all employees of a large Italian bank, with branches in every region of the country, and with a century-long tradition of activity at the heart of the Italian financial system. Our data cover all employees who worked at the firm during a three-year stretch from 1993 through 1995. For this analysis, we include only those workers who work full time and are continuously on payroll for the entire three-year stretch. The dataset provides information on the exact dates of each absence from the workplace. Our analysis focuses exclusively on absences due to illness.³ We therefore exclude all employees who took maternity leave at any point during this period.⁴ This provides a sample of 16,208 workers, of which we focus on the 14,857 who have at least one illness-related absence during the three years observed. Because of good working conditions and high salaries, attrition among the employees of this bank is negligible; only 12 workers quit and 4 were fired during this 3-year period.⁵ The descriptive statistics in Appendix Table A1 indicate that among this sub-sample of workers with at least one illness-related absence, there are 2,965 women and 11,892 men. Females are younger and slightly more educated, but have significantly more sick-days. They are also paid on average 20 percent less and are heavily under-represented

³Under Italian law, workers can take an almost unlimited number of paid sick days. In theory workers need a medical certificate if their absence extends beyond three days, but such a certificate is easily obtained. Workers are also subject to the possibility of a medical control at home, yet this control can only occur at previously specified times of the day.

⁴We also exclude the 166 top managers, of whom only two are women.

⁵An additional 4 workers died and 261 retired.

in the managerial ranks.⁶

If the menstrual cycle systematically affects female absenteeism, we should see that sick leaves of pre-menopausal women display a cycle of approximately 28 days. To investigate this hypothesis, we begin with three pieces of graphical evidence. Figure 1 shows the gender difference in the distribution of days between consecutive absences from work due to illness. In particular, the figure shows the gender difference in the distribution of number of days between the beginning of each absence, for spells that are 50 or fewer days apart. Note the spike at 27 and 28 days, indicating that the probability that consecutive spells are roughly 28 days apart is higher for women than for men. Although the graph is somewhat noisy, there are no other obvious peaks.

One limitation of this figure is that it may miss some menstrual-related absences. For instance, suppose that a woman experiences menstrual episodes precisely every 28 days, but is also absent for other reasons in between. By using only consecutive absences, Figure 1 will miss the cyclicity of some menstrual-related absences. To account for this, the top left panel in Figure 2 repeats this exercise, now including all possible pairs of absences. Again this figure shows that the probability that any two episodes are 28 days apart is higher for women than for men. The remaining panels in Figure 2 show that the spike at 28 days is driven primarily by younger workers, disappearing with age. The top-right panel, which includes only workers under 45, displays a marked difference at 28 days. This difference is less evident in the bottom-left panel, which includes workers 45 to 55, and disappears in the bottom-right panel, which includes only workers 55 or older. This pattern is consistent with the timing of menopause.⁷

An alternative way to look at cycles in absenteeism is to estimate hazard rates. Starting from the first day of a given absence spell, the top panel in Figure 3 plots Kaplan-Meier estimates of the hazard of a second absence, by gender and age, for the following fifty days. The left panel is for workers 45 or younger, the right for those over 45.⁸ Three features of these figures warrant comment. First, the hazard is almost always higher for women,

⁶Given that the firm is a bank, blue collar workers are a small minority, and this is especially the case for females. This dataset was also used by Ichino and Maggi (2000), Ichino, Polo and Rettore (2003) and Ichino and Riphahn (2004).

⁷The medical literature indicates that although many women experience menopause between 45 and 55, the age at onset varies greatly.

⁸For computational simplicity, in this Figure we focus on the first two consecutive absence episodes.

mirroring their higher overall absence rates. As mentioned in the Introduction, this pattern is common among Western countries. Second, consistent with Figure 2, the spike at 28 days is more pronounced for women under 45 than for similarly-aged men. This fact is more readily apparent in the bottom panels, which plot the female-male difference in hazards. In comparison, there is no clear spike at day 28 for those over 45, regardless of gender.

The third factor evident from Figure 3 is that both males and females have spikes at durations equal to 7, or multiples of 7. This pattern is in part driven by the “Monday morning” effect, common in many countries.⁹ For both genders, Monday is by far the most common day for the start of a sick spell: 33% and 35% of female and male absences begin on Monday, respectively. By comparison, the fraction of absences that begin on other days of the week ranges from 11% to 21% for females, and 12% to 19% for males. As a result, regardless of gender, an interval of 7 days, or multiples of 7 – including, of course, 28 – is the most common length between two consecutive absences from work. This implies that the 7-day periodicity creates a confounding element with respect to the pattern potentially induced by the menstrual cycle. Thus Figure 3 highlights the necessity to control appropriately for this confounding effect when testing for the existence of a 28-day pattern of female absenteeism.¹⁰

2.2 Parametric Hazard Estimates

Figures 1 to 3 are consistent with the hypothesis that menstrual episodes increase the risk of 28-day cyclical absences for pre-menopausal women. In this subsection we use a parametric model to test the statistical significance of this finding, controlling for the 7-day periodicity of overall absenteeism seen above and for other possible confounding factors.

While in typical applications of duration models the shape of the baseline hazard is of primary interest, here the main focus is on a specific interaction between the effect of time and the effect of gender, independent of the baseline. For this reason, we base our analysis on the partial-likelihood approach proposed by Cox (1972). Looking at two consecutive absences, we specify the hazard of the second as:

$$h(t, X_{it}, \Psi) = \lambda(t)e^{\alpha + \beta F_i + \gamma M_{it} F_i + \delta S_{it} F_i + \theta Z_{it}} \quad (1)$$

⁹For example, see Card and McCall (1996) for US evidence.

¹⁰The Monday morning effect explains some but not all of the 7 day cycle. The remaining portion is due to the fact that family and other non-work commitments often have weekly periodicity. For example, activities like own and children’s sporting events, concerts, or visits to health clinics, are all likely to repeatedly fall on the same day of each week.

where the index t represents distance in days from the previous absence; $X_{it} = (F_i, M_{it}, S_{it}, Z_i)$, $\Psi = (\alpha, \beta, \gamma, \delta)$, $\lambda(t)$ is the baseline hazard; $F_i = 1$ indicates that worker i is female; M_{it} is a dummy variables taking a value of 1 if time t is 28 ± 3 ; S_{it} is a dummy taking value of 1 if t is 7 or multiples; and Z is a vector of covariates.¹¹

The parameter β captures the overall difference in absenteeism for women relative to men. The main parameter of interest is γ . A positive estimate would indicate that females have a higher hazard of being absent from work 28 ± 3 days after their previous absence, regardless of their baseline. One important advantage of the parametric model is that it controls for the confounding time pattern induced by the 7-day periodicity of absences. If this confounding pattern is identical for males and females it will be captured by the baseline hazard. The interaction $S_{it}F_{it}$ allows for the possibility that the 7-day periodicity differs between genders; the parameter δ captures the extent to which this pattern is more or less pronounced for women.

Note that the definition of M_{it} allows for some variation in the length of a given menstrual cycle. The reason is that it is unlikely that the menstrual cycle has duration exactly equal to 28 days for all women. Menstrual cycles vary enormously, both across women, and across months for a given woman. Cycles from 25 to 31 days are not considered unusual in the medical literature. For example, Creinin, Keverline and Meyn (2004) report that for 46% of their subjects the length of the cycle can vary by 7 days or more. For 20% the length of the cycle can vary by 14 days or more. Figures 1 and 2 discussed above are consistent with this notion.

Table 2 presents estimates of the parameters e^β and e^γ . The estimated coefficients are reported in the form of hazard ratios (with the t-statistics in parentheses). For example, in the first panel, $e^\beta = 1.58$ indicates that the hazard of an absence from work is on average 58 percent higher for women than men. In addition to this higher overall risk, females experience

¹¹If we order the completed durations from the lowest to the highest ($t_1 < t_2 < \dots < t_N$ where N is the number of workers) the conditional probability that worker j concludes a spell at t_j , given that $N - j$ workers could have concluded their spell at the same time is by $\frac{h(t, X_{jt}, \Psi)}{\sum_{i=j}^N h(t, X_{it}, \Psi)} = \frac{e^{\alpha + \beta F_j + \gamma M_{jt} F_j + \delta S_{jt} F_j + \theta Z_i}}{\sum_{i=j}^N e^{\alpha + \beta F_i + \gamma M_{it} F_i + \delta S_{it} F_i + \theta Z_i}}$. This is also the contribution to the likelihood for the worker with the j th shortest duration. Note that the baseline hazard $\lambda(t)$ cancels out and does not need to be estimated. Censored observations appear in the denominator of the contribution of each observation, but do not enter at the numerator with a contribution of their own. As far as ‘‘ties’’ are concerned, i.e. units concluding the spell in the same measured time interval, we rely on the standard method consisting of including a different contribution to the likelihood for each tied observation, using the same denominator for each. This denominator includes all the tied observations.

an additional 29 percent relative increase in the hazard when $t = 28 \pm 3$. ($e^\gamma = 1.29$, statistically significantly different from 1). By comparison, the estimates of e^δ are never significantly different from one, suggesting that the 7-day periodicity is similar across genders.

In the second panel of Table 2, we divide the sample by age. Here we test the hypothesis that female absenteeism for those under 45 displays a 28-day cycle, while female absenteeism for those over 45 does not. The gender effect for overall absenteeism is similar in these two groups, indicating that the higher absenteeism of women relative to men does not disappear with age (column 1). By contrast, the incidence of the 28-day pattern is clearly different between younger and older workers (column 2). In particular, the coefficient for those under 45 is large and statistically significant, while the coefficient for those over 45 is indistinguishable from one. The hazard generated by the approximate 28-day cycle of menstruation creates an additional 44% difference in absenteeism between young women and men, while it creates no difference between older workers. To check that the results are not driven by the 7-day periodicity, we have also estimated similar models restricting the sample to those spells that do not begin on a Monday. We obtain estimates very similar to the ones in Table 2.

In the bottom panel of Table 2, we calculate the same hazards when controlling for age, years of schooling, marital status, number of children, managerial occupation, seniority and dummies for the weekday in which the spell began. Results are similar to those above. For females under 45, the hazard of an absence increases by 43 percent with respect to males at cycles of 28 ± 3 days, while again we find no such effect for those over 45.

To allow for irregular cycles, our base specification defines M_{it} as taking a value of 1 for $t = 28 \pm 3$. We also experimented with varying this interval around the 28-day mark, using plus or minus 2, 1 and 0 days. Our estimates are robust to these alternatives. In particular, columns 3 and 4 of Table 2 show these results using the most restrictive definition of M_{it} , now set equal to 1 only for $t = 28$, which captures only those cycles that are exactly 28 days long. The results do not change significantly. Consistent with the hypothesis that the menstrual cycle is indeed the driving force behind the 28-day cyclical absenteeism, the point estimates of e^γ for women under 45 are slightly higher.

The estimates in Table 2 are obtained by focusing on a cycle centered at 28 days. Yet we can perform the same exercise for other lengths of time, equivalent to letting the data tell us the correct periodicity of the cycle of female absenteeism. Finding a significant effect for

cycles that are not centered around 28 days would cast some doubt on the interpretation of our results. Appendix Table A2 reports these results when we pretend that the menstrual cycle exerts its effect in periods different than the biologically-driven one. Restricting the analysis to females younger than 45, each row comes from a different regression, in which we change the periodicity of the cycle. The estimate of e^γ for the window centered at 28 – which corresponds to the estimate in the second row and second column of Table 2 – is the largest and the most precise. It also has the highest log likelihood. The other coefficients and corresponding log likelihoods decline monotonically as we move further away from 28. The intervals centered at 21 and 35 days – both of which may capture a non-trivial number of short and long menstrual cycles – are marginally significant. None of the remaining intervals are. These results confirm the visual impression obtained from Figures 1 and 2. Finally, we have also estimated the same models for workers older than 45 (results not shown). As expected, none of the coefficients are different from zero, with the exception of the interval at 21 days, which is only marginally so.¹²

2.3 How many days of work are lost in connection to the menstrual cycle?

Above we established that there is a statistically significant increase in the hazard of an absence for young females every 28 days. We now want to know whether this phenomenon is not only statistically significant, but also quantitatively relevant. In this section we therefore estimate the number of days of work lost per year because of the 28-day cycle of menstruation, and we report the extent to which our estimates match the existing medical literature.

We focus on workers who are 45 years old or younger, and we consider the distance between all pairs of short absences from work. In particular, we define an absence as short if it lasts 3 or fewer days, and we call two absences cyclical if they are both short, and they are between 26 and 30 days apart, or multiples thereof. The assumption is that menstrual symptoms are unlikely to induce long absences.

¹²We tried to obtain data on pill use to see whether the incidence of the 28 days cycle is different in areas where pill use is more widespread. Unfortunately, available data on pill use in Italy are not disaggregated geographically. But even if the data were available, it is not exactly clear what to expect, since the effect of pill use on the 28 days cycle is ambiguous. On one hand, pill use reduces the pain caused by menstrual cramps. On the other hand, pill makes the cycle more regular, and therefore more likely to be measured in the data.

Based on this definition, we compute the total number of cyclical absence pairs for each worker in our sample. We then normalize this by the number of pairs of all short absences from work experienced by that employee. We therefore obtain an index ranging from 0 to 1, that represents the worker-specific fraction of short absences that has an (approximate) cycle of 28 days.¹³

It is important to realize that even for men this indicator may be larger than zero, although it should be on average smaller than for women. There are two reasons for this. First, and most importantly, male absenteeism has a 7-day periodicity. Because 28 is a multiple of 7, men have a certain number of absences that appear to be characterized by a 28-day cycle, even if they clearly do not suffer menstrual symptoms. Second, men as well as women are likely to experience a certain number of 28-day cyclical absences just by chance. For example, it is possible that some workers experience two illnesses 28 days apart that have nothing to do with menstruation. For our purposes, the key implication of the 7-day cycle and of the possibility of false positives is that we can only identify the average number of absences induced by menstrual episodes as the *difference* between women and men in the measured number of absences with a 28-day pattern.

Figure 4 shows that, as expected, women have a much larger fraction of absences with a 28-day pattern. The figure plots, by gender, the cumulative distribution of the fraction of cyclical absences. It is apparent that the distribution for women stochastically dominates the distribution for men.¹⁴

¹³Our results are robust to alternative definitions of the cycle. First, we obtain similar results when two absences are defined cyclical if they are exactly 28 days apart, between 27 and 29 apart, or between 25 and 31 days apart (or multiples thereof). Second, our results do not change if instead of considering all possible multiples, we only include the first 5 multiples. In other words, our results are mostly driven by pairs of absences that are 5 cycles or less apart. This makes sense, because it is unlikely that the menstrual cycle is so regular that menstrual episodes that are several months apart are still aligned on a 28 day cycle. Third, our results are also robust to changes in the definition of a short absence. For example, they remain essentially unchanged when we define an absence as short if it lasts 2 days or less, or 4 days or less.

¹⁴One potential concern is that the number of false positives is larger for women than men because women have higher absenteeism. As a consequence, the estimated difference in the number of cyclical absences may overestimate the true number of menstrual episodes. To get a sense of whether this problem is empirically relevant, we have calculated the theoretical number of false positives for men and women. In particular, given 365 days, we have simulated the timing of absence episodes under the null of no cyclical absenteeism. We assume that each episode is i.i.d., and that the timing of each absence is uniformly distributed over the course of the year. We use different distributions for men and women, so that the number of episodes for men and women is equal to their respective averages reported in Table 3. Using 1000 repetitions, we find that the difference in the number of false positive is negligible. (3.3% of all pairs of absences is a false positive for men, while the corresponding number for women is 3.4%.) We conclude that this problem is unlikely to affect our estimates in any significant way.

To obtain an estimate of the number of days of cyclical absences for each worker, we multiply the worker-specific fraction of cyclical absences by the worker-specific number of short absences. Table 3 quantifies the gender difference in total and cyclical absenteeism. The first row indicates that men in the sample have on average 8.2 days of absence each year, while women have 12.9 days. The resulting gender difference in absenteeism is therefore 4.6 days. The second row shows our estimates of the number of days of cyclical absences. The unconditional gender difference is now 1.4 days (column 2). This difference is our best guess of the effect of menstrual episodes on absenteeism for the average woman. Based on this difference, we conclude that about 30% of the gender difference in days of absenteeism is due to menstrual symptoms ($1.4/4.6 = 0.3$). The gender difference conditional on age and education is 1.5 days (columns 4 and 5).

Rows 3 and 4 show similar figures for the number of episodes of absenteeism. Here the importance of the menstrual cycle is even more evident. For example, column 3 indicates that women have on average 1.5 more absence spells than men. The corresponding figure for cyclical absences is 1.1. This implies that 73% of the gender difference in episodes of absenteeism may be due to menstrual symptoms ($1.1/1.5 = 0.73$).

These average gender differences mask large variation in the distribution of days of cyclical absences. Table 4 therefore shows this distribution by gender. The distribution for women is clearly shifted to the right of that for men. The fraction of men and women for whom the number of estimated cyclical absences is 0 is 55% and 29%, respectively. By contrast, the fraction of men and women for whom the number of estimated cyclical absences is 2 or more days is instead 25% and 49%.

Medical Evidence. One way to assess the plausibility of our estimates is to compare them with the existing medical literature. Various studies in this literature report estimates suggesting that as many as 75% to 90% of premenopausal women regularly experience some form of mild premenstrual symptoms.¹⁵ A smaller fraction of women typically meet all the criteria for the clinical definition of pre-menstrual syndrome, or for its more severe version,

¹⁵See, among others, Johnson (1987), Deuster et al. (1999), Sternfeld et al. (2002) and Chawla et al. (2002). The premenstrual syndrome (PMS) is typically defined in the medical literature as “a cluster of physical and emotional symptoms that appear on a regular basis before the onset of menstrual bleeding. Symptoms include bloating, breast pain, ankle swelling, a sense of increase in body weight, irritability, aggressiveness, depression, lethargy and food cravings.” (Deuster et al. 1999).

the “premenstrual dysphoric disorder” (PMDD). Much of the existing research is focused on the possible association between PMS and behavioral outcomes such as suicide, psychiatric hospitalization, criminal activity, accidents and work performance (Johnson 1987). From our point of view, the frequency, regularity and severity of premenstrual symptoms is relevant inasmuch as it interferes with the normal working life of affected females.

In a recent study specifically aimed at measuring the “economic burden” of the premenstrual syndrome, Chawla et al. (2002) provide the most comprehensive medical evidence to date. A representative sample of 1,194 California women aged 21 to 45 was asked to provide prospective daily symptom ratings and information on health care use and work productivity for two menstrual cycles. The estimates in Chawla et al. (2002) of the number of days of activity lost due to the menstrual cycle are remarkably similar to our estimates. Specifically, their estimates imply that the average woman in their sample experienced about 1.7 cut-down days in a year because of physical symptoms associated with the menstrual cycle.¹⁶ Our estimates in Table 3, based on the same age range, indicate that the average woman in our sample experienced about 1.4 days of absence. Remarkably, it is not just the mean that it is similar in the two samples, but the distribution of menstrual episodes across women also appears to be similar. Although the two samples are not homogenous because they come from different countries and involve a different occupational mix, we conclude that our estimates are not implausible when compared with the best existing medical evidence.

2.4 Real pain or shirking?

The estimates presented so far are consistent with the hypothesis that the menstrual cycle increases the hazard of an absence from work for pre-menopausal females. This does not necessarily mean, however, that the reason for this increase in the hazard is the physical symptoms caused by menstruation. It is possible that taking a day off from work in association with one’s menstrual cycle is still a matter of choice, and that menstruation simply offers women a socially acceptable occasion to shirk.

¹⁶They report that 17.3% of their sample had “severe” symptoms. For 5% of the sample the severity was so high to originate a PMDD diagnosis. While even the most severe symptoms induced little bed time per menstrual cycle, at least 1.1 days were cut down from work and other usual activities by the 17.3% of women who experienced severe symptoms (1.3 days for PMDD women). Since their figures are based on two menstrual cycles, the implied number of cut-down days is obtained as follows: $(1.1 * 0.17 + 1.3 * 0.05) * (365 / (28 * 2)) = 1.7$.

In this subsection and in the following Section we try to investigate whether variation across women in the documented 28-day cycle of absenteeism reflects shirking or, alternatively, whether it reflects an exogenous and largely unavoidable health shock. In particular, in this subsection, we test whether the 28-day cycle is less pronounced for workers for whom the cost of shirking is higher. We present three pieces of evidence. First, we show that the 28-day cycle is equally pronounced for managers and clerks. Second, we show that the cycle is no less pronounced for those workers who are up for a promotion. And third, we show that the cycle is no more pronounced in bank branches where the overall level of absenteeism is higher.

Taken as a whole, we interpret this evidence as suggesting that variation across women in the incidence of absenteeism related to the menstrual cycle does not simply reflect variation in the propensity to shirk. In the next Section, we present more evidence consistent with this hypothesis by empirically testing the predictions of a model of wage determination where workers differ in their propensity to shirk (Section 3).

Managers vs Clerks. We begin in Table 5 by looking at the incidence of the 28 day cycle by occupational level. Columns 1 and 2 are based on managers; columns 3 and 4 on clerks. The idea is that the former should be less likely to shirk than the latter. Finding that the 28-day cycle is limited to lower rank employees and does not affect women in higher positions would suggest that variation in the 28-day cycle is driven more by heterogeneity in worker quality than by real physical symptoms.¹⁷

Consistent with the notion that managers should be less likely to shirk, managers in our sample have a much lower propensity to be absent. Managers take on average 5.2 sick days per year, compared to 10.6 days for clerks. Moreover, the gender differential in absenteeism for young managers is slightly lower than that for young clerks. This is seen by comparing the estimate of e^β in columns 1 and 3 of Table 5.

Although managers have substantially lower absenteeism than clerks overall, their incidence of the 28-day cycle is not in fact lower. If anything, the incidence is *higher* for managers, as documented by the fact that estimates of e^γ in column 2 are significantly

¹⁷We define managers and clerks broadly. Specifically, the category "managers" include workers in hierarchical levels 7 to 13, therefore including both upper management ("dirigenti") and supervisors ("quadri"). The category "clerks" include workers in hierarchical levels 1 to 6. This group includes mostly white collar workers, but also includes a small number of manual occupations, like janitors.

higher than the corresponding estimates for clerks in column 4. As one might expect, this is true for younger, but not older, workers. For managers, the hazard of an absence due to illness for females relative to males more than doubles during the days at risk of a menstrual cycle. By contrast, the increase for clerks is 28%. Estimates in the bottom panel, where we control for worker characteristics, are very similar.

Before and After a Promotion. In Table 6 we focus on workers who received a promotion during this three-year period, and test whether the incidence of the 28-day cycle is different before and after the promotion. The idea is that the signaling cost of an absence in the months leading up to a promotion is higher than in the months immediately following.¹⁸ Finding that absences associated with menstrual symptoms are more likely to occur after a promotion than before would again suggest that shirking could be an important determinant of the observed 28-day cycle of absenteeism.

We find that in the year after a promotion, workers have slightly higher overall absenteeism than in the year before. This is true for both men and women.¹⁹ This is not surprising, since workers have strong incentives to minimize their absenteeism in the months leading up to a promotion decision. Remarkably, however, even if the overall level of absenteeism is lower before a promotion, the incidence of 28-day cyclical absences is not lower beforehand. Table 6 shows that, if anything, the incidence is in fact *higher* before a promotion than afterwards; the coefficients for workers under 45 shown in the top panel are 3.87 before a promotion and 2.87 after. Although the sample is small, the estimates remain significantly different from 1. By contrast, the estimates for workers over 45 are not statistically significant.²⁰

In interpreting these estimates, it is important to realize that we only observe the selected sample of promotions that actually occur. We do not observe those cases where an employee was considered for a promotion, and did not receive it. Therefore, our estimates may not

¹⁸We include one year before the promotion and one year after the promotion. We only consider *merit promotions*, i.e. promotions based on performance. We do not include promotions based on seniority, because those promotions do not depend on performance, but occur automatically based on a set schedule.

¹⁹The average number of sick days before and after the promotion is 3.35 and 4.17 for men; 5.07 and 5.70 for women. These numbers are lower than the average number of sick days for the whole sample, presumably because workers who experience a merit promotion are less likely to shirk.

²⁰One possible explanation for why the effect is larger before the promotion is that stress is known to heighten PMS.

generalize. In particular, we have no way of telling how the incidence of the 28-day cyclical absences may change after a promotion decision for workers who failed to obtain a promotion. The problem is, of course, that one might expect that workers who did obtain a promotion have a lower propensity to shirk than workers who failed to obtain a promotion.

Although other explanations are certainly possible, Tables 5 and 6 are consistent with the notion that the periodicity in absenteeism is not lower when the cost of such absenteeism is high. Taken together, these pieces of evidence seem to indicate that physical symptoms increase during the days at risk of a menstrual cycle, and that women affected have limited freedom to decide whether or not to go to work on those days.

Work Environment. A related question is whether the hazard attributable to menstrual effects changes with the work environment. In particular, does the observed 28-day periodicity differ in bank branches where average absenteeism is high, relative to branches where average absenteeism is low? Similarly, does it vary depending on how many women are employed in the branch? To answer these questions, we modify model 1 to allow for the triple interaction of the female indicator, the indicator for the 28-day cycle, and an environment variable, E_{it} :

$$h(t, X_{it}, \Psi) = \lambda(t)e^{\alpha + \beta F_i + \gamma M_{it} F_i + \phi M_{it} F_i E_{it} + \psi E_{it} + \delta S_{it} F_i + \theta Z_{it}} \quad (2)$$

In column 1 of Table 7, the variable E_{it} is the fraction of females in the relevant branch. In column 2, E_{it} is a dummy equal to one if the employee works in the South. Southern branches are characterized by significantly higher absenteeism than branches in other regions. This is likely due to the fact that Southern Italy is characterized by a work culture more tolerant of absenteeism (Maggi and Ichino, 2000).

Estimates in column 1 of Table 7 indicate that a prevalently female work environment is associated with substantially higher overall absenteeism. Going from an all-male to an all-female branch is associated with a 52 percent increase in the hazard of an absence.²¹ Yet, a prevalently female work environment is not associated with an increase in the effect of

²¹Of course this finding should not necessarily be interpreted causally. For example, it is possible that females are prevalently segregated in less important branches of the firm in which the chances of a promotion are lower. Every worker in these branches would have a smaller incentive to reduce absenteeism, but this would not be due to the prevalence of females per se. Alternatively, a higher prevalence of females might generate an environment where absence episodes due to family duties are more socially acceptable. In this case, the effect of the proportion of females on absenteeism could be interpreted causally.

menstrual cycles. While females continue to have a higher hazard relative to men, even after controlling for the fraction of women in the branch, the increase in the hazard at 28 days does not depend on this fraction: the interaction coefficient e^ϕ is not significant.

Column 2 shows that the hazard of an absence is approximately 30 percent higher for employees working in branches located in Southern Italy (see column 2). Even in the presence of this large overall difference, however, the incidence of absenteeism in a 28-day cycle does not appear to be different in the North than in the South. Females in the South experience the same hazard of an absence during the days at risk of a menstrual cycle: the interaction coefficient e^ϕ is not significant.²²

Overall, Table 7 indicates that the association between the menstrual cycle and absenteeism does not change between environments with high or low rates of overall absenteeism. This finding is consistent with the idea that our measure of the menstrual cycle does not depend on social norms that define the socially acceptable level of absenteeism.

3 Gender Differences in the Relationship Between Absenteeism and Earnings

In the previous Section we have argued that absences showing a 28-day cycle explain a significant fraction of the male-female absenteeism gap. In this and the next sections we are interested in quantifying the effect of this source of absenteeism on the gender gap in earnings. We begin in this section by presenting a simple model that clarifies how the relationship between absenteeism, worker quality, and earnings may differ for men and women. The model provides a set of testable implications that we bring to the data. The predictions of the model become useful in the next Section, when we use variation across workers in the incidence of the 28-day cycle to quantify the effect of the menstrual cycle on the earnings gender gap. In particular, we use the predictions of the model to evaluate the validity of the identifying assumption needed for the counterfactual calculation.

²²We have also estimated models where E_{it} is simply the overall level of absenteeism in a branch. Consistent with what we find in Table 7, the coefficient on the triple interaction is not statistically significant. We do not report these estimates, however, since they are not immediately interpretable because of the “reflection problem” discussed, among others, by Manski (1993).

3.1 A Simple Model

The idea of the model is simple. If employers cannot directly observe individual productivity, they might use observable worker characteristics, including absenteeism, to predict productivity and set wages.²³ In particular, we assume that employers set wages according to a simple model of statistical discrimination, weighting their gender-specific priors and the observed signal. The key insight is that absenteeism is a noisier signal of shirking attitudes for females than for males. The reason is that inasmuch as menstrual-related absences are not a signal of shirking, signal extraction based on absenteeism is more informative about shirking for males than for females. The implication is that we should expect the relationship between earnings and absenteeism to be more negative for males than for females. Furthermore, this gender difference in slope should decline with seniority, since the informational content of absenteeism declines as employers learn more about their workers. This last point has been made in a different context by Farber and Gibbons (1996) and Altonji and Pierret (2001).

We begin by considering a model without workers' effort. Below, we show that our main result does not change when we allow workers to endogenously choose their effort level to reduce absenteeism, knowing that this decision will affect their future wages.

Absenteeism has two types of effects on workers earnings. First, there is a mechanical loss of output. Workers who are often absent produce less and in the long run earnings will adjust. Second, there might be a signaling effect in cases where employers have imperfect information on worker quality. For notational simplicity, we focus on this second type of cost. Including the mechanical effect of absenteeism does not change our results. Assume that the productivity per unit of working time of employee i is given by

$$Y_i = c - S_i \tag{3}$$

where c is a constant and S_i is the individual propensity to shirk. We think of S_i as a measure of worker i 's permanent quality; workers with large S_i are those with permanently higher propensities to shirk. The firm, however, observes neither S_i nor Y_i . The firm instead

²³We modify the career concerns model proposed by Holmstrom (1999). See also Aigner and Cain (1977). In a different context, Milgrom and Oster (1987) use asymmetric information to explain gender wage differences. For an empirical example of the link between absenteeism and shirking see Skogman Thoursie (2004). See also Hotz (1984).

observes only absenteeism, X_{it} , in period t , and pays the wage²⁴

$$W_{it} = E(Y_i|X_{it}) \quad (4)$$

We think of male absenteeism as the sum of non-menstrual health shocks and the propensity to shirk. Female absenteeism is caused by these two factors, plus menstrual-related absences. Menstrual-related absences are assumed to be exogenous health shocks that have nothing to do with shirking. In particular, we assume that

$$X_{it} = S_i + \mu H_{it} \quad (5)$$

where μH_{it} are health shocks and S_i and H_{it} are independent. Although employers can observe X_{it} , they do not know whether an absence is caused by a real health shock (H_{it}) or by shirking (S_i). In other words, the worker has no way to credibly signal which absences are caused by real illness. The effect of menstrual episodes is captured by the loading factor μ . To capture the idea that females have more health shocks than males because of the menstrual cycle, we assume that $\mu = 1$ for males and $\mu > 1$ for females. Note that the term X_{it} can represent either total absenteeism or cyclical absenteeism. Because the focus of this paper is on menstrual-related absenteeism, in our empirical application X_{it} will represent absenteeism with cycle 28 days.

We also assume that

$$S_i \sim N(\omega, \frac{1}{p}) \quad (6)$$

$$H_{it} \sim N(\eta, \frac{1}{q}) \quad (7)$$

where the parameters ω, η, p, q and μ are known to everyone. These assumptions amount to stating that because of the menstrual cycle, the distribution of cyclical absenteeism for females has both a higher mean and – more importantly – a higher variance than the distribution for males. This assumption appears realistic. In our sample, women’s cyclical (and total) absences have a significantly higher mean and variance than men’s.

We first consider what happens in period 1. Employers use the Normal Learning Model to predict which workers are productive and which workers are shirkers, based on the level

²⁴Because productivity and earnings are measured in *units of working time*, workers are paid only for the time when they are on the job. For this reason, the cost of absenteeism in the model is purely its signaling value.

of observed absenteeism:

$$\begin{aligned} E(S_i|X_{i1}) &= E(S_i + \mu\eta|X_{i1}) - \mu\eta \\ &= \frac{p}{p + \frac{q}{\mu^2}}(\omega - \frac{q\eta}{\mu}) + \frac{\frac{q}{\mu^2}}{p + \frac{q}{\mu^2}}X_{i1} \end{aligned} \quad (8)$$

This updating rule simply says that the employer's best guess of the unobserved propensity to shirk of worker i is a precision-weighted average of the data (X_{i1}) and the prior ($\omega - \frac{q\eta}{\mu}$).²⁵ As a result, given equation 4, the wage paid by the firm in period 1 is

$$W_{i1} = c - E(S_i|X_{i1}) = \alpha - \beta X_{i1} \quad (9)$$

where α is a constant and the slope in the relationship between earnings and absenteeism is

$$\beta = \frac{\frac{q}{\mu^2}}{p + \frac{q}{\mu^2}} \quad (10)$$

It is easy to see that

$$\beta_{male} > \beta_{female} \quad (11)$$

The intuition behind the difference in slope coefficients, $\beta_{male} > \beta_{female}$, is that absenteeism is a noisier signal of productivity for females than for males, and therefore observed absenteeism has a larger effect on employer's priors for men. Thus an absence episode is associated with a smaller earnings loss for women than men.

Having shown that the relationship between earnings and absenteeism is initially more negative for men than women, we can now determine how these relationships evolve over time, as employers learn more about each worker's quality, S_i . Learning here is captured by iteration of the Normal Learning Model. Iterating the Normal Learning Equation, we can see that after t periods – as more information becomes available – Equation 8 becomes:

$$E(S_i|X_{i1}, \dots, X_{it}) = \frac{p}{p + t\frac{q}{\mu^2}}(\omega - \frac{q}{\mu}\eta) + \frac{\frac{q}{\mu^2}}{p + t\frac{q}{\mu^2}} \sum_{s=1}^t X_{is} \quad (12)$$

This equation implies that with the passage of time, the precision of the prior on the individual propensity to shirk improves for both genders, until S_i becomes fully known in the limit.

²⁵Unlike the most commonly used models of statistical discrimination, in our model the crucial difference between genders is not the difference in the mean of the prior, but in the variance of the prior.

The wage offer in period t can therefore be expressed as a function of the worker-specific average absenteeism up to period t , $\bar{X}_{it} = \frac{1}{t} \sum_{s=1}^t X_{is}$:

$$W_{it} = \alpha_t - \beta_t \bar{X}_{it} \quad (13)$$

where $\alpha_t = c - \frac{p}{p+t-\frac{q}{\mu^2}}(\omega - \frac{q}{\mu}\eta)$ and

$$\beta_t = \frac{\frac{q}{\mu^2}}{\frac{p}{t} + \frac{q}{\mu^2}} \quad (14)$$

The key implication is that as t goes to infinity, the slope β_t becomes -1 , *irrespective of gender*. The intuition is that when the information on S_i available to the employer increases, the fact that observed absenteeism is a more noisy measure of shirking for females becomes increasingly less relevant. With perfect information (i.e. when t is equal to infinity), the signal becomes completely irrelevant, and any gender difference in the relationship between earnings and absenteeism disappears.²⁶

One reasonable question is whether workers really have no control over health-related absenteeism. For example, one might think that for a given health shock, a worker can reduce her absenteeism by exerting effort and showing up for work even if she does not feel very well. Our framework can be generalized to include effort decisions and career concerns. Endogenizing effort as in Holmstrom (1999) allows workers to decide how much effort to exert knowing that this decision will affect their future wage via the employer signal extraction process.

In the Appendix, we show that this generalization does not change the basic result of our model. When effort is considered explicitly, women anticipate that their observed absenteeism is a more noisy measure of shirking propensity and that an additional absence episode is less costly for them than for men. As a consequence, women have a lower incentive to exert effort. Notably, the β coefficient in equation 11 remains unchanged. The slope in the relationship between earnings and absenteeism remains negative for both genders and

²⁶Note that the slope does not go to zero with perfect information, because workers with high absenteeism have, by assumption, a higher propensity to shirk. An alternative way of seeing the same result is the following. The wage offer in period t can be expressed as a function of the worker-specific observations on absenteeism X_{is} up to period t : $W_{it} = \alpha_t + \sum_{s=1}^t \beta_s X_{is}$ where $\beta_s = \frac{\frac{q}{\mu^2}}{p+s-\frac{q}{\mu^2}}$. The key implication is that as s goes to infinity, the slope β_s becomes 0, *irrespective of gender*. With perfect information, the additional signal at time t becomes irrelevant.

smaller in absolute value for women.²⁷ The main effect of introducing effort is that the gender gap in absenteeism widens, since women’s equilibrium effort is lower.

3.2 Empirical Tests of the Model

In this subsection we take the model to the data. While the model does not distinguish between total absenteeism and cyclical absenteeism, in this paper we are interested in the latter. Specifically, in the next Section we seek to identify a counterfactual gender gap in earnings in the absence of *cyclical* absenteeism. There, we use the predictions of the model on the relationship between worker quality and cyclical absenteeism to assess the validity of our key identification assumption. For this reason, our empirical tests in this Section focus on cyclical absenteeism. In practice it is reasonable to assume that the employer can observe not only total yearly absenteeism for each worker, but also the timing of each absence. This is realistic, since the firm keeps track of the exact date and duration of each absence. Indeed, the firm collects—and presumably uses—the same data that we use. If the employer can observe the timing of each absence, it can identify which absences have a 28 days cycle.

The three testable implications of the model can be summarized as follows. A first implication on the relationship between earnings and cyclical absenteeism follows immediately from equation 11, which indicates that the cost of a day’s absence is lower for women than for men.

Proposition 1. *If menstrual episodes do not reflect shirking, in a regression of earnings on cyclical absenteeism the coefficient is negative for both genders, but the coefficient is smaller in absolute value for females than for males.*

A second prediction on how the relationship between earnings and absenteeism varies over time follows from equation 14. In taking equation 14 to the data, we face an important data limitation. Because our sample includes only three years of data, we do not have enough time variation to test equation 14 using longitudinal data. Instead, we use cross-sectional differences across workers with different seniority levels. The idea is that the employer has more information on workers with more seniority than workers with less seniority.

²⁷This does not contradict the career concern motive described by Holmstrom (1999). It remains true that females, like males, exert more effort at the beginning of their careers, but on average, over the entire career, they will exert less effort than males in reducing absenteeism.

Proposition 2. *In a regression of earnings on cyclical absenteeism, if the slope coefficient on cyclical absenteeism differs initially by gender, it will become more similar across genders as seniority increases.*

Finally, an additional implication involves the relationship between worker quality and absenteeism.

Proposition 3. *In a regression of measures of worker quality on cyclical absenteeism, the coefficient is negative for both genders, but the coefficient is smaller in absolute value for females than for males. Moreover, any gender difference in the slope coefficient on absenteeism will remain constant as seniority increases.*

We think of worker quality as the inverse of the propensity to shirk, S_i . The first part of Proposition 3 is easily derived using equation 5.²⁸ The second part of Proposition 3 derives from the fact that under our assumptions, X_{it} is a stationary variable and its correlation with the time invariant propensity to shirk should not change over time.

We now provide empirical tests for these three predictions of the model. We note that the predictions from the model do not necessarily involve causality, since they are simply equilibrium outcomes.

Earnings and Careers. The entry in the first column in the top panel of Table 8 shows that the unconditional earnings gap in our sample is -20%.²⁹ This is consistent with the earnings gap observed in the US for white-collar workers. When we control for a quadratic in age and the number of non-cyclical absences in column 2, the earnings gap declines to

²⁸Given that $X_{it} = S_i + \mu H_{it}$ an hypothetical regression of S_{it} on X_{it} would yield a slope coefficient equal to

$$\frac{\text{cov}(S_i, X_{it})}{\text{var}(X_{it})} = \frac{\text{var}(S_i)}{\text{var}(S_i) + \mu^2 \text{var}(H_{it})} = \frac{\frac{1}{p}}{\frac{1}{p} + \frac{\mu^2}{q}} \quad (15)$$

because S_i and H_{it} are orthogonal. Since μ is larger for females than males, equation 15 implies a steeper positive slope for males if the dependent variable is the propensity to shirk S_i (and therefore a steeper negative slope for males if the dependent variable is worker quality.) Note that in this case the OLS coefficient in equation 15 is identical to the parameter β of the Normal Learning Model in equation 10. However, OLS and the Normal Learning Model are in general not the same thing. OLS applies to a situation in which both the dependent variable and the independent variable (or at least their proxies) are observed. In the Normal Learning Model, the conditional expectation of S_i given X_{it} can be obtained even if S_i unobserved. This is possible because of functional form assumptions.

²⁹In this and all the remaining Tables, we use only workers who are 45 or younger.

-13.5%.

Column 3 is a direct test of Proposition 1. Log earnings are regressed on a dummy for female, the yearly number of cyclical absences, and the interaction of female and cyclical absences.³⁰ The estimates are consistent with Proposition 1. Increases in cyclical absences are associated with declines in earnings, both for males and females. But the decline is significantly less steep for females than males. An additional day of cyclical absences costs male workers about 2.5%. The cost for female workers is only 1.5%.³¹ Since in this specification we include workers of any seniority, we interpret the estimated coefficients as an average across all seniority levels. Below we let the coefficients differ based on seniority.

In the two remaining panels, we look at the relationship between cyclical absenteeism and careers. In this firm, there is a tight correspondence between earnings and occupational rank, and there is limited variation in earnings within an occupational level. The main way in which workers obtain a raise is by being promoted to a higher level. For this reason, the findings in the top panel are qualitatively similar to those in the middle and bottom panel, where the dependent variable is occupational rank.

Specifically, in the middle panel, the dependent variable is a dummy for whether the worker is ever promoted to manager. Women are 18% less likely to be promoted to a management position (column 1), or 11% when controls are included (column 2). Consistent with Proposition 1, when we include measures of absenteeism interacted with gender, the probability of promotion to management declines with absenteeism for both men and women, but the decline is significantly more marked for men (column 3). In the bottom panel 3, the dependent variable is a linear measure of occupation. This model assumes that the distance between occupational levels is the same at each promotion step. In this data there are 13 occupational categories. For example, the dependent variable for executives is equal to 13, for supervisors it is 8, for tellers 6, for junior tellers 5, and for manual occupations it is 1.

³⁰Earnings in this firm are paid monthly, and are not mechanically adjusted for the number of days of absences in that month. So there is not a immediate month-to-month relationship between earnings and cyclical absences. However, earnings are presumably adjusted in the longer run to reflect absenteeism.

³¹Note that these effects are larger than daily earnings for several reasons. First, these effects reflect not only the direct loss of output caused by the absence, but also the signaling value of avoiding absences. As the model indicates, a worker who is often absent is more likely to be considered a shirker by the employer. Second, there are fixed costs (capital, insurance, etc.) paid by the firm, irrespective of whether the worker is on the job or absent. Third, the cost of an unplanned absence to an employer includes the disruption that it causes, and therefore is likely to be larger than daily earnings. We will come back to these points in Section 4.1.

The mean (std deviation) of the dependent variable is 6.1 (2.2). Again, the estimates shown in Table 8 are consistent with Proposition 1.³²

Effect of Seniority. We now turn to a test of Proposition 2. The lack of longitudinal data leads us to use cross-sectional differences across workers with different seniority levels. Since very few workers in this firm are fired or quit (Section 2.1), selective attrition is not a significant concern. The specification in Table 9 generalizes the one in column 3 of Table 8 by including the triple interaction of female, cyclical absences, and years of seniority, as well as including each of the main effects and their pairwise interactions. The coefficient of interest is the one on the triple interaction. Proposition 2 predicts that the gender difference in the earnings-absenteeism relationship decreases with seniority, because the employer learns more about a worker's true propensity to shirk. Therefore the prediction is that the coefficient on the triple interaction should be negative.

As in Table 8, we find that an increase in cyclical absences is associated with a significant decline in earnings, both for men and women. When seniority is low, the negative slope in this relationship is steeper (more negative) for men. For example, when seniority is 0, the slope in this relationship is -.032 for men and only -.019 for women. More importantly, this gender difference in the slope in the relationship between earnings and cyclical absenteeism declines with seniority. Consistent with Proposition 2, the coefficient on the triple interaction is negative and statistically significant: -.007 (.0003). To better see the effect of seniority on the difference in slope, Figure 5 shows predicted earnings for men and women, relative to cyclical absences, for those with 0 years of seniority (left panel) and 15 years of seniority (right panel). The figure shows that when workers first join the firm, men and women have different slopes. After 15 years, when the employer has learned about individual productivity, this difference in slope disappears.³³

³²Alternative interpretations are of course possible. For example, assume that workers' tasks differ in how easily they can be performed by a substitute worker in case of absence. Specifically, assume that the cost of an absence of a worker whose task can easily be performed by a substitute is lower than the cost of an absence of a worker whose task cannot be performed by a substitute. If women are more likely to be absent, profit-maximizing management should be more likely to assign tasks that can easily be performed by substitutes to women than men. This would explain the lower cost of a day of absence for women. This explanation is likely to be more relevant in firms with large heterogeneity in tasks. Although we do not have data on tasks, we suspect that in this firm's tasks are fairly homogenous. For instance, the tasks performed by clerks in most branches of this bank are quite standardized.

³³Note that the effect of seniority on earnings is not immediately apparent, because in the regression we

This finding is important because it lends further credibility to the notion that when a worker first join this firm, the employer can not observe true individual productivity, but learns about it only over time. If the employer could observe true individual productivity from day one, we would not see the change in slope that we uncover in Figure 5. In other words, this finding appears to support our assumption that employers have initially imperfect information and that they use absenteeism as a signal.

Worker Quality. Finally, we turn to the relationship between worker quality and absenteeism described in Proposition 3. The first part of Proposition 3 indicates that if we could observe worker quality, we should see a steeper decline in quality for men than for women as absenteeism increases. Obviously we have no good measure for S_i , but only some imperfect proxies. For this reason we stress that this evidence is to be considered only as suggestive.

In column 1 of Table 10 we use schooling. The results are consistent with Proposition 3. For men, increases in cyclical absenteeism are associated with a steep decline in schooling. The coefficient is -0.104, indicating that each additional day of cyclical absence is associated with a decline in schooling of one tenth of a year. For women, there is effectively *no* relationship. The coefficient is $-0.01 = -0.104 + 0.094$, and not statistically different from zero.³⁴

Similarly, in column 2 the dependent variable is a dummy equal to 1 if the worker is involved in any misconduct episodes in the three years observed.³⁵ Consistent with Proposition 3, workers with more cyclical absences are more likely to have been sanctioned, but this is significantly less true for women than for men.

In column 3 the dependent variable is the number of days of vacation taken. While all workers have a right to the same amount of vacation (five weeks per year), there is substantial variability in the actual number of days taken. The assumption here is that workers who

control for age. In the Figure, we fix age and non-cyclical absences to be equal to the age and non-cyclical absences of the average worker in the sample.

³⁴Of our four proxies, schooling is the one that is the most predetermined, and therefore the one that is closest to the notion of worker quality S_i . The other three proxies are relatively less predetermined.

³⁵These are episodes where worker misconduct is recorded and punished by the personnel office. The punishments vary in terms of severity, from verbal reproach to the ultimate level, firing. In this column, we exclude managers, because they are not subject to misconduct sanction.

take only part of their allotted vacation days are more driven and career-oriented than others. The point estimates indicate that higher absenteeism is associated with more days of vacation for men, but not for women. However, the standard errors are too large to draw firm conclusions. Similarly, if we instead use days of strike (column 4), the signs are as predicted, but the estimates are again too small and imprecise to allow interpretation.

Finally, the second part of Proposition 3 implies that in a regression of exogenous measures of worker quality on absenteeism, any gender difference in the slope coefficient on absenteeism should remain constant as seniority increases. This stands in contrast to Proposition 2. In this sense, a test of Proposition 3 can be considered a specification test of the evidence on Proposition 2. Finding that gender differences in the relationship between exogenous measures of worker quality and absenteeism vary over time in the same way that gender differences in the relationship between earnings and absenteeism do would cast some doubt on the interpretation of our test of Proposition 2. We estimated specifications similar to the one in Table 9, using as dependent variables the four indicators of workers quality that we use in Table 10. In all cases, we found that the coefficient on the triple interaction between cyclical absence, seniority and gender is statistically insignificant.³⁶

4 How Much of the Gender Gap in Earnings is Explained by the Menstrual Cycle?

In most countries women earn less than men. In our sample of Italian bank workers, women earn about 13.5% less than the men, conditional on observables (Table 8, column 2). The magnitude of this earnings difference is similar to that observed in representative samples from other countries. For example, in the US the conditional gender gap for white-collar workers in this same age range is approximately -15%. In other European countries it is about -17%. In this Section, we try to determine how much of the observed gender gap in earnings is explained by the absenteeism generated by the menstrual cycle.³⁷ There are two ways to answer this question.

³⁶The coefficient (std. error) on the triple interaction for schooling is .0005 (.003); for misconduct, .0004 (.0004); for vacation, .0001 (.004); for strike, -.0009 (.0009).

³⁷The literature on the causes of gender differentials in labor market outcomes is immense, and can not be summarized here. For a recent survey of the literature, see Altonji and Blank (1999). For recent trends in the gender gap see Blau and Kahn (forthcoming). For an international comparisons see Blau and Kahn (2003).

4.1 Direct Effect

We begin by measuring the direct effect of menstrual-related absenteeism by calculating the value of work time lost due to the menstrual cycle:

(Days of work lost due to cycle \times Women's average daily earnings) / Gender gap in earnings

Our estimates in Table 3 suggest that the 28-day cycle is associated with 1.5 days of additional absenteeism for the average woman. Given approximately 214 working days per year (excluding week-ends, holidays, and vacations), and given that average earnings for women and men are 25,020 and 29,034 Euros, respectively, 4.4% of the earnings gap can be explained by the direct effect of this absenteeism on earnings: $[1.5 \times (25020/214)]/4014 = 4.4\%$.

However this direct effect is only part of the total effect of the menstrual cycle on the earnings gap. The reason is that the effect of a day of absence on earnings is arguably larger than daily earnings, for several reasons. First, the calculation above does not reflect the signaling value of avoiding absences. As the model indicates, absences may be used by employers to distinguish between shirkers and non-shirkers. As a consequence, the cost of a day of absence for a worker should include both the value of lost output as well as the cost of sending a bad signal.

Second, this estimate does not reflect the fixed costs (capital, insurance, etc.) paid by the firm, irrespective of whether the worker is on the job or absent. Third, this estimate does not reflect the lost productivity due to menstrual symptoms when the worker is on the job. It is possible that there are instances when a female worker experiences menstrual symptoms that lower her productivity, but in which the pain is just below her threshold to trigger an absence. Medical studies confirm that women's on-the-job productivity declines substantially as a consequence of menstrual symptoms.³⁸ Fourth, in most white-collar jobs the cost to the employer of the disruption caused by an unplanned day of absence is surely

³⁸For example, in a clinical study, Chawla et al. (2002) estimate that women with severe PMS symptoms experience decreases in productivity of 48.2%—64.4% for women with the more severe PMDD—relative to the women with minimal symptoms. The decline in productivity was measured using productivity scores computed according to the Endicott Work Productivity Scale (Endicott, 1997) and time diaries. All the differences are statistically significant at the 1% level. Self-assessed productivity declines were between 13.8% and 22.7%. Chen et al. (2005) find evidence of lower performance for women during menstruations in a laboratory experiment.

more than the daily earnings of that person.

4.2 Total Effect

For all these reasons, the estimate of the effect of cyclical absenteeism based on daily earnings shown above is likely to be a lower bound. We now adopt an alternative approach to determine the total effect of menstrual cycle on the gender gap in earnings. We divide workers into groups according to the number of their cyclical absences. We then re-weight the groups using a counterfactual distribution, based on the observed male distribution of cyclical absenteeism.

Specifically, we can write men’s earnings as $Y_m = \pi_{1m}Y_{1m} + \pi_{2m}Y_{2m} + \pi_{3m}Y_{3m}$, where Y_{1m} , Y_{2m} , and Y_{3m} are the average earnings of men in the groups with a low, medium and high number of cyclical absences, respectively. π_{1m} , π_{2m} , and π_{3m} are the fractions of men in each group. Similarly, we can write women’s earnings as $Y_f = \pi_{1f}Y_{1f} + \pi_{2f}Y_{2f} + \pi_{3f}Y_{3f}$. Empirically, we define the groups so that they have equal size. Workers in group 1 have no cyclical absences, while those in groups 2 and 3 have on average of 1.1 and 4.5 days, respectively. Because of the menstrual cycle, men are over-represented in the lowest group and women are over-represented in the highest. In particular, 49 percent of the men are in group 1, compared to only 22 percent of the women. In contrast, only 28 percent of the men are in group 3, versus 52 percent of the women.

The observed difference in earnings between females and males is simply:

$$Y_f - Y_m = (\pi_{1f}Y_{1f} - \pi_{1m}Y_{1m}) + (\pi_{2f}Y_{2f} - \pi_{2m}Y_{2m}) + (\pi_{3f}Y_{3f} - \pi_{3m}Y_{3m}) \quad (16)$$

What would be the earnings gap if a women did not suffer from menstrual symptoms? We estimate the counterfactual earnings gap by assigning to everyone the distribution across groups for men:

$$\tilde{Y}_f - \tilde{Y}_m = \pi_{1m}(Y_{1f} - Y_{1m}) + \pi_{2m}(Y_{2f} - Y_{2m}) + \pi_{3m}(Y_{3f} - Y_{3m}) \quad (17)$$

Equation 17 provides a counterfactual earnings gap by moving some women from group 3 to groups 1 and 2, and some women from group 2 to group 1, so that the distribution of men and women in cyclical absences is equalized.

To empirically quantify the gender difference in earnings for the three groups, we estimate the following equation:

$$\log Y_i = \beta_1 + \beta_2 C_{2i} + \beta_3 C_{3i} + \gamma_1 C_{1i} F_i + \gamma_2 C_{2i} F_i + \gamma_3 C_{3i} F_i + \mu X_i + e_i \quad (18)$$

where F_i is an indicator for females; C_{ji} is an indicator for the j group of the cyclical absences distribution ($j = 1, 2$ or 3); and X_i controls for non-cyclical absences and age. The parameters of interests are the γ 's, which are our estimates of the gender earnings gap for each group: $\gamma_1 = (Y_{1f} - Y_{1m})$; $\gamma_2 = (Y_{2f} - Y_{2m})$; and $\gamma_3 = (Y_{3f} - Y_{3m})$. This approach relaxes the linearity assumption implicit in the models in Table 8, allowing the relationship between earnings and absenteeism to be non-linear.

This strategy provides the valid counterfactual gender gap under two identifying assumptions:

Assumption 1. The menstrual cycle is the only reason for a difference between men and women in the number of days of absences with a 28-day cycle. Note that this assumption requires that the menstrual cycle is the only cause of the male-female difference in *cyclical* absences. It does not require that the menstrual cycle is the only cause of the male-female difference in the *total* number of days absent from work. Obviously there are many reasons why the total absenteeism distribution is different for women and men.

Assumption 2. In estimating equation 18, the female-male difference in unobservables is the same for all three groups, or at least it does not decline with cyclical absences. Specifically, we assume that

$$\begin{aligned} E(e_i|F, j = 3) - E(e_i|M, j = 3) &\geq E(e_i|F, j = 2) - E(e_i|M, j = 2) \\ &\geq E(e_i|F, j = 1) - E(e_i|M, j = 1) \end{aligned} \quad (19)$$

This assumption is about how gender *differences* in unobservables vary across groups. It is not about how *levels* of unobservables vary across groups. For example, this assumption does not rule out the likely possibility that ability declines with absenteeism, so that women in group 1 have higher average ability than women in group 3. However, it does rule out the possibility that the difference in ability between women and men declines with absenteeism, so that the difference in ability between women and men in group 1 is larger than the corresponding difference in group 3.

Note that if the difference in unobservables is the same across groups, estimates of the effect of menstrual cycle on the gender gap are unbiased. If instead the female-male difference in unobservables increases with cyclical absences, then estimates of the effect of menstrual cycle on the gender gap are a *lower bound* of the true effect. The intuition is that a positive correlation between gender differences in unobservables and cyclical absences would lead us to underestimate the effect of cyclical absences on women’s earnings, and therefore the estimated effect of menstrual cycle on women’s earnings.

Is this assumption plausible? Proposition 3 predicts that the gender difference in worker quality should *increase* with days of cyclical absences. In other words, the model predicts that the average female-male difference in worker quality is smallest in group 1 and largest in group 3. If the model is correct, our identification assumption is not violated and our estimates should be interpreted as a lower bound of the true effect of menstrual cycle. This prediction of the model is corroborated by the evidence in Table 10, which confirms that for some measures of worker quality, the female-male difference increases with cyclical absenteeism.

Findings. Table 11 shows estimates of the β ’s and γ ’s in equation 18. Entries in column 1 show that earnings decline for both men and women as we move from group 1 to groups 2 and 3. The earnings difference between groups 1 and 3 is statistically significant for both men and women, as indicated by the two tests at the bottom of the table. Consistent with Proposition 1 and with Table 8, the decline is steeper for men than for women. When we re-weight the γ ’s using the male distribution across the three groups (equation 17), we estimate that the counterfactual gender gap is -11.9% (see the bottom of the Table). This should be compared with the observed gender gap from Table 8, column 2, namely -13.5%. We conclude that if women did not experience 28-day cyclical absenteeism, the earnings difference between females and males would be 1.6 percentage points—or 11.8 percent—lower than the observed difference.

Three comments are worth making. First, we can compare this estimate, 11.8 percent, with the estimate obtained above of the direct cost of absenteeism, 4.4 percent. The latter figure is an estimate only of the direct cost of absenteeism, i.e. the value of work time lost due to menstrual symptoms. The 11.8 percent figure includes the direct effect, as well as the signaling value of absenteeism, the value of any fixed costs, the value of lost productivity on

the job and the cost of disruption in case of unplanned absences. The comparison suggests that the direct cost represents only about a third of the total cost of absenteeism for a worker.

Second, the estimates in Table 11 are simply a more general version of the linear models in Table 8. The main advantage is that these estimates allow for a non linear relationship between earnings and cyclical absenteeism. Estimates based on linear models are slightly smaller, although not very different.³⁹

Third, the way to interpret this counterfactual gap is as the earnings gap that we would observe if we eliminated menstrual symptoms for a given woman, *holding fixed the incidence of menstrual symptoms for all other women*, and then averaged these counterfactual female earnings for all women. By holding fixed the incidence of menstrual symptoms of all other women, we are effectively holding fixed the gender difference in the cost of an absence. This counterfactual gap is conceptually different from the gap that we would observe if all women did not suffer menstrual symptoms, since presumably in this case the price of an absence faced by women would change.

Finally, columns 2 and 3 repeat the same exercise using an indicator of career progression as the dependent variable. The effect of 28-day cyclical absenteeism on the career gender gap is 13.5% or 24.4%, depending on whether the outcome variable is a management dummy, or the measure of occupational level.

5 Conclusions

In most countries women take more sick days than men. We argue that an important cause of this gender difference may be the menstrual cycle. Absenteeism of those women in our sample who are 45 or younger displays a systematic pattern with a cycle of approximately 28 days. Absenteeism of women who are 45 or older shows no such cyclical pattern. Overall, a third of the gender gap in days of absence, and two-thirds of the gender gap in the number of absence spells, appears to be due to the menstrual cycle. The incidence of

³⁹The effect of menstrual-related absenteeism based on linear models can be calculated using the following formula: $[(\text{Days of work lost due to cycle} \times \text{Cost of a day of cyclical absence for women}) / \text{Gender gap in earnings}]$ Note that this is similar to the formula used in Section 4.1. The only difference is that we have substituted “daily earnings” with “cost of a day of cyclical absence”. Our estimates in Table 8 suggest that a day of cyclical absence costs women 1.5% of earnings. Given that women earn on average 25,020 Euros, the formula implies that about 9.3% of the earnings gap can be explained by the direct effect of this absenteeism on earnings: $[1.5 \times (25,020 * 0.015)] / 4014 = 9.3\%$. We thank Claudia Goldin for suggesting this calculation.

cyclical absenteeism remains significant even for those workers who one would expect to be less likely to shirk, namely managers and workers who are in line for a promotion.

What is the effect on women's earnings and careers of this additional absenteeism? Using a simple model, we argue that an important component of the cost of an absence comes from its signaling value. If employers cannot directly observe productivity, they may set wages using workers' observable characteristics, including their propensity to be absent. But because of menstrual related absences, absenteeism is a noisier measure of worker quality for females than for males. Consistent with the prediction of the model, we find that earnings are a declining function of absences, and that this decline is steeper for men than for women. Thus while females have more cyclical absences than males because of the menstrual cycle, a cyclical absence costs more for men than for women. This difference in slope disappears with seniority, however, as employers acquire more information on workers' true productivity.

We estimate how much of the observed gender gap in earnings and careers can be attributed to the additional absenteeism induced by the menstrual cycle. The gender gap in earnings in our sample is -13.5%. Using a simple re-weighting scheme, we calculate that if the average woman did not suffer menstrual symptoms (while all other women did), the gender gap would decline to -11.9%. In other words, the gender gap in earnings would be 11.8 percent lower. A similar calculation shows that the gender gap in the probability of being promoted to manager would be 13.5% lower. These figures are likely to be lower bounds, because the decline in worker quality associated with increases in absenteeism should be weaker for women than men.

We stress that our findings are based on data from only one firm and their external validity is unclear. On the other hand, our estimates of the incidence of menstrual related absenteeism match remarkably well medical estimates based on a representative sample of Californian women. Women in the two samples come from different countries, have different occupations, are subject to different labor market institutions and incentives, and yet, they seem to have similar cyclical absenteeism. Clearly, more research is needed to verify if the same relationship between cyclical absenteeism and earnings is observed in other contexts.

References

- Aigner and Cain, "Statistical Theories of Discrimination in labor markets", *Industrial and labor Relations* (1977) , 175-187
- Altonji, Joseph and Rebecca Blank, "Gender and Race in the Labor Market" in *Handbook of Labor Economics*, Volume 3C. Orley C. Ashenfelter and David Card, eds. New York, NY: Elsevier Science Press(1999)
- Altonji, J. and C. Pierret, "Employer Learning and Statistical Discrimination", *Quarterly Journal of Economics*, 116(1) (February 2001): 313-350.
- American College of Obstetricians and Gynecologists (ACOG), (2000), "Premenstrual Syndrome", in *ACOG Practice Bulletin* n. 15.
- Barmby T. A., Orme C.D. and Treble J. C. (1991), "Worker Absenteeism: An Analysis Using Microdata" *Economic Journal*, 101, 214-229.
- Bertrand, Marianne and Kevin Hallock, "The Gender Gap in Top Corporate Jobs," in *Industrial and Labor Relations Review* (2001)
- Blank, Rebecca "Simultaneously Modeling the Supply of Weeks and Hours of Work Among Female Household Heads" in *Journal of Labor Economics*. Vol 6:2. p177-204 (1988).
- Blau, Francine and Lawrence Kahn, "The US Gender Pay Gap in the 1990s: Slowing Convergence," in *Industrial and Labor Relations Review*, forthcoming.
- Blau, Francine and Lawrence Kahn "Understanding International Differences in the Gender Pay Gap," *Journal of Labor Economics* (2003)
- Bridges S. and K. Mumford (2000), "Absenteeism in the UK: A Comparison Across Gender", University of York, WP 2000/12.
- Card D. and B. P. McCall (1996), "Is Workers' Compensation Covering Uninsured Medical Costs? Evidence from the Monday Effect", *Industrial and Labor Relations Review* 49, 4, 690-706.
- Chawla Anita, Barbara Sternfeld, Ralh Swindle, Stacey Long and Sean Kennedy, (2002) "Premenstrual Dysphoric Disorder. Is There an Economic Burden of Illness?", *Medical Care*, 40, 11, 1101-1112.
- Chen Y., Katuscak P. and Ozdenoren E (2005), "Why Can't a Women Bid More Like a Men", mimeo, Stanford University.
- Creinin, M. Keverline S. and Meyn L. (2004) "How regular is regular? An analysis of menstrual cycle regularity", *Contraception*, 70(4):289-92.
- Currie, Janet and Rosemary Hyson, "Is the Impact of Health Shocks Cushioned by Socio-economic Status" in *American Economic Review*, May 1999, 89-2, 245-250.

- Currie, Janet, "Male Jobs, Female Jobs, and Gender Gaps in Benefits Coverage", in *Research in Labor Economics*, v14, 171-210 (1995)
- Deuster Patricia A., Tilahun Adera and Jeannette South-Paul, (1999), "Biological, Social and Behavioural Factors Associated with Premenstrual Syndrome", *Arch. Fam. Med.*, 8, 122-128.
- Endicott J., Nee J., (1997), "Assessment measures for clinical studies. Endicott Work Productivity Scale: A new measure to assess treatment effects ", *Psychopharmacological Bulletin*, 33, 13-6.
- Farber, Henry and Robert Gibbons (1996), "Learning and Wage Dynamics," *Quarterly Journal of Economics* 111 (1996): 1007-47.
- Holmstrom B. (1999), "Managerial Incentive Problems - A Dynamic Perspective" in: *Essays in Economics and Management in Honor of Lars Wahlbeck*. Helsinki: Swedish School of Economics, reprinted in: *Review of Economic Studies* 66(1), 169-182.
- Hotz, Joe "Statistical Approaches to Modeling Absenteeism" in *Absenteeism*, P. Goodman et al. eds., San Francisco: Jossey-Bass, Inc., 1984, pp. 158-193.
- Hotz, Joe and R. Miller, "An Empirical Analysis of Life Cycle Fertility and Female Labor Supply," in *Econometrica*, Vol. 57, No. 1, January 1988.
- Ichino A. and G. Maggi (2000), "Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm", *Quarterly Journal of Economics* 115(3), 1057-1090.
- Ichino, A., M. Polo, and E. Rettore (2003), "Are Judges Biased by Labor Market Conditions?", *European Economic Review* 47(5), 913-44.
- Ichino A. and R. Riphahn (2004), "The Effect of Employment Protection on Worker Effort. A Comparison of Absenteeism During and After Probation", *Journal of the European Economic Association*, 3:1, 120-143.
- Johnson S. R., (1987), "The Epidemiology and Social Impact of Premenstrual Symptoms", *Clin Obstet Gynecol.*, 30, 367-376.
- Kiefer N. M. (1988), "Economic Duration Data and Hazard Functions", *Journal of Economic Literature* XXVI, June, 646-679.
- Leigh J. P. (1983), "Sex Differences in Absenteeism", *Industrial Relations* 22:3, Summer, 349-361.
- Manski Charles F., Identification of Endogenous Social Effects: The Reflection Problem, *Review of Economic Studies*, 60, 1993, pp. 531-542.
- Milgrom, Paul and Oster, Sharon, "Job Discrimination, Market Forces, and the Invisibility Hypothesis," *The Quarterly Journal of Economics*, 1987.

- Oyer, Paul, "The Structure of Wages and Internal Mobility", *AER Papers and Proceedings*, 94, May 2004, 212-216.
- Paringer L. (1983), "Women and Absenteeism: Health or Economics", *American Economic Review* 73:2, May, 123-127.
- Skogman Thoursie Peter (2004), "Reporting sick: are sporting events contagious?", *Journal of Applied Econometrics* 19, 809-823.
- Sternfeld Barbara, Ralh Swindle, Anita Chawla, Stacey Long and Sean Kennedy, (2002) "Severity of Premenstrual Symptoms in a Health Maintenance Organization Population", *Obstet Gynecol.*, 99, 1014-1024.
- VandenHeuvel A. and Wooden M. (1995), "Do Explanations of Absenteeism Differ for Man and Women", *Human Relations* 48:11, 1309-1329.
- Vistnes J. P. (1997), "Gender Differences in Days Lost from Work due to Illness", *Industrial and Labor Relations Review* 50:2, January, 304-322.

Appendix: Absenteeism with Endogenous Effort

In this Appendix, we generalize the simple model of Section 3 to include effort decisions. The idea is that workers know that absenteeism is used by employers to predict productivity and set wages. Therefore, for a given health shock, workers may exert effort to reduce the negative signal of an absence. Employers are aware of this, and set wages accordingly.

The timing is the following. First, employers offer an optimal take-it-or-leave-it wage schedule, as a function of absenteeism. Workers observe their cost of effort, shirking propensity and health shocks and choose effort optimally. This determines observed absenteeism of workers and in turn their wage. We retain most of the structure of the previous model and focus on period 1. We modify equation 5 as

$$X_{i1} = S_i - e_{i1} + \mu H_{i1} \quad (20)$$

where e_{i1} represent the effort that worker i can exert to reduce absenteeism. The employer, anticipating an optimal choice of effort on the part of the worker, offers the following wage schedule:

$$E(S_i|X_{i1}) = \frac{p}{p + \frac{q}{\mu^2}}(\omega - \frac{q\eta}{\mu p}) + \frac{\frac{q}{\mu^2}}{p + \frac{q}{\mu^2}}(X_{i1} + e_{i1}^*) \quad (21)$$

where e_{i1}^* is the optimal effort choice of the worker, to be defined below. Assume that exerting effort is costly and workers maximize $W_{i1} - \frac{\theta_i e_{i1}^2}{2}$, where the parameter θ_i characterizes the cost of effort of worker i .⁴⁰ Optimal effort is therefore

$$e_{i1}^* = \frac{1}{\theta_i} \frac{\frac{q}{\mu^2}}{p + \frac{q}{\mu^2}} \quad (22)$$

As a result the equilibrium wage is

$$\begin{aligned} W_{i1} &= c - \frac{p}{p + \frac{q}{\mu^2}}(\omega - \frac{q\eta}{\mu p}) - \frac{1}{\theta_i} \left(\frac{\frac{q}{\mu^2}}{p + \frac{q}{\mu^2}} \right)^2 - \frac{\frac{q}{\mu^2}}{p + \frac{q}{\mu^2}} X_{i1} \\ &= \bar{\alpha} - \beta X_{i1} \end{aligned} \quad (23)$$

We draw two main conclusions. First, compared with equation 9, the slope coefficient β is unchanged. In particular, it remains steeper for men than women. Second, even if the distribution of cost of effort parameter θ_i is the same for men and women, the intercept $\bar{\alpha}$ differs from the intercept α in equation 9. In particular, the difference between the intercept for men and women is now larger because women have a lower incentive to exert effort.

⁴⁰For simplicity, we assume that θ_i is orthogonal to S_i and H_{i1} .

Table 1: Gender Differences in Days of Absence in a Year, by Country

	All workers		Unmarried No Children	
	(1)	(2)	(3)	(4)
Europe	6.67** (0.52)	7.65** (0.60)	2.12** (0.80)	2.78** (0.88)
Usa	3.07** (0.23)	3.09** (0.43)	1.09** (0.49)	2.01** (0.88)
Canada	5.22** (0.09)	5.19** (0.11)	0.31 (0.17)	1.13** (0.20)
Our Sample	4.66** (0.32)	5.04** (0.33)	2.76** (0.53)	3.70** (0.54)
Controls	N	Y	N	Y

Notes: Each entry is the gender difference (females - males) in the number of days of absence from work in a year. Samples include full time workers not on maternity leave. Controls in columns 2 and 4 include age, education level dummies, occupational qualification dummies. Controls in column 2 also include the number of children and marital status, and country specific dummies for the European sample. Standard errors in parentheses with $p < 0.05 = *$, $p < 0.01 = **$. The top row uses data from the European Community Household Panel (N=38,229). Row 2 uses data from the PSID (N = 11,735). Row 3 uses data from the Canadian Labor Force Survey (N=575,243).

Table 2: Hazard of an absence for females relative to males and the risk of a menstrual cycle.

	Cycle has periodicity 28 ± 3 days		Cycle has periodicity 28 days	
	e^β	e^γ	e^β	e^γ
<u>All Ages</u>	1.58	1.29	1.59	1.15
	(19.94)	(2.78)	(20.84)	(0.75)
<u>By Age Group</u>				
Under 45	1.55	1.44	1.57	1.49
	(17.50)	(3.51)	(18.47)	(1.89)
Above 45	1.58	0.97	1.58	0.35
	(7.70)	(-0.10)	(7.82)	(-1.67)
<u>By Age Group, With Controls</u>				
Under 45, with controls	1.56	1.43	1.58	1.49
	(16.72)	(3.47)	(17.61)	(1.88)
Above 45, with controls	1.43	0.96	1.43	0.35
	(5.82)	(-0.14)	(5.89)	(-1.67)

Notes: Asymptotic t-ratios in parentheses. Entries are the Cox-Proportional Hazard ratios for the occurrence of a second absence episode with time measured from the beginning of the first absence episode after January 1, 1993, computed from the estimation of equation 1. e^β is the hazard ratio of females relative to males in a day not at risk of a menstrual cycle; e^γ is the factor by which the hazard ratio of females relative to males increases in a day at risk of a menstrual cycle. Entries in the bottom panel are obtained conditional on age, years of schooling, marital status, number of children, managerial occupation, seniority and dummies for the weekday in which the spell begins. Sample sizes are 14857 (row 1), 10793 (row 2), 4064 (row 3), 10793 (row 4), and 4064 (row 5).

Table 3: Absenteeism, by Type and Gender

	Men	Women		Difference	
	(1)	(2)	Unconditional	Conditional	Conditional
	(1)	(2)	(3)	(4)	(5)
<u>Days of Illness-Related Absence</u>					
Total Number of Days in a Year	8.2	12.9	4.6	5.2	5.4
			(.3)	(.3)	(.3)
Estimated Number of Cyclical Days in a Year	1.3	2.9	1.4	1.5	1.5
			(.06)	(.06)	(.06)
<u>Episodes of Illness-Related Absence</u>					
Total Number of Episodes in a Year	2.1	3.6	1.5	1.6	1.6
			(.5)	(.5)	(.6)
Estimated Number of Cyclical Episodes in a Year	.9	2.0	1.1	1.1	1.1
			(.04)	(.04)	(.04)
Control for Age			N	Y	Y
Control for Education			N	N	Y

Notes: Standard errors in parenthesis. Sample includes workers 45 or younger.

Table 4: Distribution of Number of Days of Cyclical Absences in a Year, by Gender

Number of Days of Cyclical Absence	Frequency for Males (1)	Frequency for Females (2)
0	55	29
1	20	22
2	10	15
3	5	9
4	3	6
5	2	5
6	1	4
7	1	2
8	0	2
9	0	1
10+	1	4

Notes: Sample includes workers 45 or younger.

Table 5: Hazard of an absence for females relative to males and the risk of a menstrual cycle, for managers and clerks.

	Managers		Clerks	
	e^β	e^γ	e^β	e^γ
<u>By Age Group</u>				
Under 45	1.41 (5.80)	2.10 (2.80)	1.49 (14.03)	1.28 (2.16)
Above 45	1.46 (3.91)	0.32 (-1.53)	1.47 (5.01)	1.07 (0.28)
<u>By Age Group, with Controls</u>				
Under 45	1.45 (6.03)	2.07 (2.75)	1.56 (14.91)	1.27 (2.09)
Above 45	1.35 (2.97)	0.32 (-1.56)	1.45 (4.72)	1.07 (0.29)

Notes: Asymptotic t-ratios in parentheses. Entries are Cox-Proportional Hazard ratios for the occurrence of a second absence episode with time measured from the beginning of the first absence episode after January 1, 1993, computed from the estimation of equation 1. e^β is the hazard ratio of females relative to males in a day not at risk of a menstrual cycle; e^γ is the factor by which the hazard ratio of females relative to males increases in a day at risk of a menstrual cycle. In the bottom panel, controls include age, years of schooling, marital status, number of children, managerial occupation, seniority and dummies for the weekday in which the spell begins. The category “Managers” include workers in hierarchical levels 7 to 12 (it therefore includes lower level management). The category “Clerks” include workers in hierarchical levels 1 to 6. Days at risk of menstrual cycle is defined as $t = 28 \pm 3$. Sample sizes in columns 1 and 2 are 3302 (row 1), 2351 (row 2), 3302 (row 3), and 2351 (row 4). Sample sizes in columns 3 and 4 are 7491 (row 1), 1713 (row 2), 7491 (row 3), and 1713 (row 4).

Table 6: Hazard of an absence for females relative to males and the risk of a menstrual cycle, before and after a promotion

	Before promotion		After promotion	
	e^β	e^γ	e^β	e^γ
<u>By Age Group</u>				
Under 45	1.52 (3.48)	3.87 (2.91)	1.75 (4.34)	2.87 (2.03)
Above 45	1.70 (1.55)	0.00 (0.00)	1.87 (1.66)	0.00 (0.00)
<u>By Age Group, with Controls</u>				
Under 45	1.46 (2.81)	3.77 (2.85)	1.64 (3.42)	2.89 (2.04)
Above 45	1.54 (1.17)	0.00 (0.00)	1.94 (1.60)	0.00 (0.00)

Notes: Asymptotic t-ratios in parentheses. Entries are Cox-Proportional Hazard ratios for the occurrence of a second absence episode with time measured from the beginning of the first absence episode after January 1, 1993. The sample includes only workers who were received a merit promotion between 1993 and 1994. The coefficients are computed from the estimation of equation 1. e^β is the hazard ratio of females relative to males in a day not at risk of a menstrual cycle; e^γ is the factor by which the hazard ratio of females relative to males increases in a day at risk of a menstrual cycle. In the bottom panel, controls include age, years of schooling, marital status, number of children, managerial occupation, seniority and dummies for the weekday in which the spell begins. Days at risk of menstrual cycle is defined as $t = 28 \pm 3$. Sample sizes in columns 1 and 2 are 523 (row 1), 207 (row 2), 523 (row 3), and 207 (row 4). Sample sizes in columns 3 and 4 are 478 (row 1), 176 (row 2), 478 (row 3), and 176 (row 4).

Table 7: Hazard of an absence for females relative to males and the risk of a menstrual cycle, by fraction of women in the branch and by region

	(1)	(2)
Females (e^β)	1.53 (15.83)	1.57 (16.93)
Females in days at risk (e^γ)	1.66 (2.69)	1.32 (2.39)
Females in days at risk \times fraction of females in branch (e^ϕ)	0.53 (-0.92)	
Fraction of females in branch (e^ψ)	1.52 (4.23)	
Females in days at risk \times South (e^ϕ)		1.27 (1.45)
South (e^ψ)		1.28 (10.53)

Notes: Asymptotic t-ratios in parentheses. Hazard ratios computed from the estimation of equation 2. Sample includes workers 45 or younger. In column 1 the indicator of work environment is the fraction of females in the worker's branch, while in column 2 it is a dummy equal to one if the employee works in the south. All models control for age, years of schooling, marital status, number of children, managerial occupation, seniority and dummies for the week-day in which the spell begins. Sample size is 10793.

Table 8: Earnings and Career Equations - Linear Models

	(1)	(2)	(3)
<u>Model 1: Earnings</u>			
Female	-.204 (.006)	-.135 (.006)	-.144 (.008)
Cyclical Absences			-.025 (.001)
Female \times cyclical absences			.010 (.002)
<u>Model 2: Promoted to Manager</u>			
Female	-.183 (.009)	-.111 (.009)	-.138 (.012)
Cyclical Absences			-.029 (.002)
Female \times cyclical absences			.017 (.003)
<u>Model 3: 13 Occupation Levels</u>			
Female	-.754 (.050)	-.216 (.047)	-.368 (.060)
Cyclical Absences			-.184 (.000)
Female \times cyclical absences			.108 (.015)
Controls for non-cyclical absences	N	Y	Y
Controls for age	N	Y	Y

Notes: Standard errors in parenthesis. In model 2 the dependent variable is a dummy equal 1 if the worker is promoted to manager or supervisor by 1995. The mean (std deviation) of the dependent variable is .24 (.43). In model 3 there are 13 occupational categories. For example, the dependent variable for executives is equal to 13, for supervisors is 8, for senior tellers is 7, for middle tellers is 6, for junior tellers is 5, for manual occupations is 1. The mean (std deviation) of the dependent variable is 6.1 (2.2). Sample includes workers 45 or younger.

Table 9: The Relationship Between Earnings and Cyclical Absences, by Gender and Firm Seniority

	Dependent Variable is Earnings (1)
Female \times Cyclical Absences \times Seniority	-.0007 (.0003)
Female \times Cyclical Absences	.013 (.002)
Female \times Seniority	.002 (.001)
Cyclical Absences \times Seniority	.001 (.0002)
Female	-.145 (.011)
Cyclical Absences	-.032 (.001)
Seniority	-.011 (.000)

Notes: Standard errors in parenthesis. Seniority is measured in years. Also included are controls for non-cyclical absences and age. Predicted earnings by gender are plotted in Figure 5. Sample includes workers 45 or younger.

Table 10: Gender Differences in Observable Indicators of Workers Quality, by Amount of Cyclical Absences

	Schooling	Misconduct	Days of Vacation Taken	Days of Strike
	(1)	(2)	(3)	(4)
Average Dep. Variable	13.1	.09	19.5	.97
Days of Cyclical Absence	-.104 (.014)	.028 (.001)	.018 (.015)	.002 (.003)
Days of Cyclical Abs. \times Female	.094 (.023)	-.018 (.002)	-.029 (.025)	-.000 (.005)
Controls	Y	Y	Y	Y

Notes: Standard errors in parenthesis. Controls include a dummy for females, a quadratic in age and dummies for number of days of non-cyclical absences. Sample includes workers 45 or younger.

Table 11: Earnings and Career Equations: Workers are Divided into Three Groups Based on the Number of Cyclical Absences

	Earnings	Manager	13 Occupations
	(1)	(2)	Levels
	(1)	(2)	(3)
Medium number of cyclical absences (β_2)	-.042 (.007)	-.037 (.010)	-.251 (.053)
High number of cyclical absences (β_3)	-.118 (.007)	-.135 (.010)	-.821 (.053)
Small number of cyclical absences \times female (γ_1)	-.131 (.012)	-.107 (.019)	-.302 (.092)
Medium number of cyclical absences \times female (γ_2)	-.118 (.012)	-.116 (.018)	-.222 (.091)
High number of cyclical absences \times female (γ_3)	-.099 (.009)	-.059 (.014)	.142 (.070)
Test coeff on High $<$ 0 (p-value)	.00	.00	.00
Test coeff on High \times Fem. $>$ coeff on Low \times Fem. (p-value)	.02	.02	.00
Observed Gender Gap (conditional)	-.135	-.111	-.216
Counterfactual Gender Gap (conditional)	-.119	-.096	-.161
Percent of the Observed Gap "Explained" by Cycle	11.8%	13.5%	25.4%

Notes: Standard errors in parenthesis. The estimated equation is equation 18. The excluded group is males with a small number of cyclical absences. All models control for the number of non-cyclical absences and for age. The observed gender gap is the coefficient on the female dummy in a regression that includes controls for the number of non-cyclical absences and for age (see column 2 in Table 8). The counterfactual gender gap is defined in equation 17. Sample includes workers 45 or younger.

Figure 1: Gender Differences in the Distribution of the Distance Between Consecutive Absence Spells.




Figure 2: Gender Differences in the Distribution of the Distance Between Absence Pairs, Using All Possible Pairs.

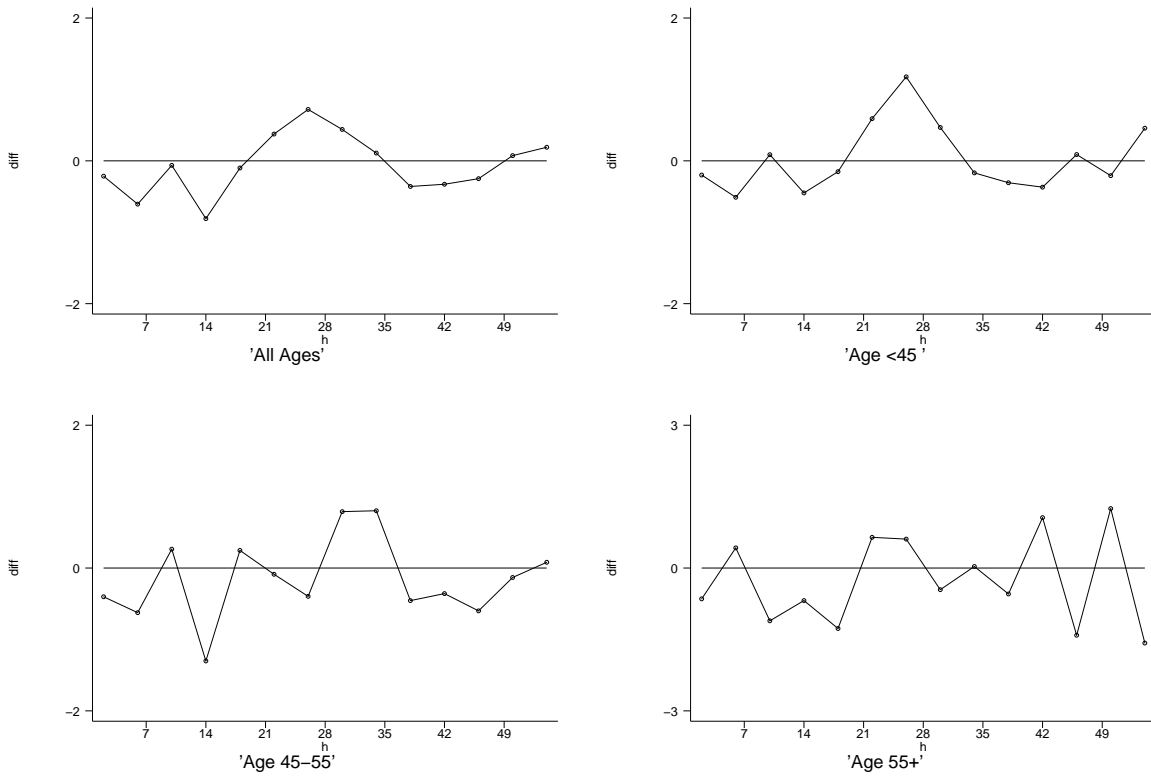
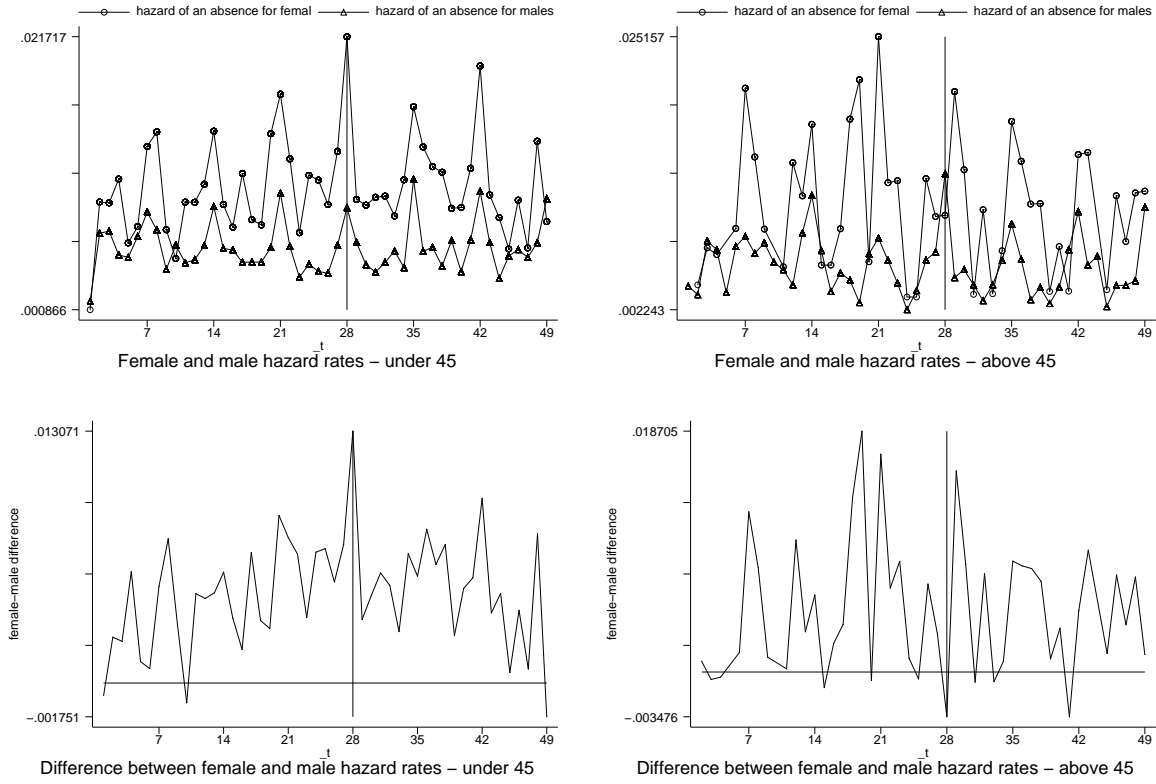



Figure 3: Hazard Rates, by Gender and Age




Note: The top panels plot the Kaplan-Meier estimates of the hazard for male and females. The bottom panels plot the difference.

Figure 4: The Distribution of the Fraction of Cyclical Absences of Women Stochastically Dominates the Distribution of Cyclical Absences of Men



Notes: the line with x is the cumulative distribution of the fraction of cyclical absences for men. The line with circles is the cumulative distribution of the fraction of cyclical absences for women. Sample includes workers 45 or younger

Figure 5: The Relationship Between Predicted Earnings and Cyclical Absences, by Gender and Firm Seniority



Note: The lines show the predicted log earnings as a function of days of cyclical absences based on estimates of the model in Table 9. The left panel is for workers with 0 years of seniority. The right panel is for workers with 15 years of seniority. Predicted earnings are for a worker of average age and average number on non-cyclical absences.

Appendix Table A1: Descriptive statistics

	Females	Males
Sick Days in a Year	12.9 (16.5)	8.2 (13.3)
Age	35.6 (7.9)	40.3 (7.8)
Years of schooling	13.3 (2.7)	13.0 (3.3)
Seniority	13.0 (7.7)	16.2 (7.9)
Yearly earnings (Euros)	25,020 (7,261)	29,034 (14,336)
Percent working in the south	25.7	28.9
Percent manager or supervisor	8.4	29.4
Percent clerk	90.7	65.9
Percent blue collar	0.9	4.6
Number of observations	2965	11892

Note: Sample includes full time workers continuously on the payroll between January 1, 1993 and December 31, 1995 who are absent at least once for illness related reasons. Workers on maternity leave are excluded.

Appendix Table A2: Placebo analysis.

Days at risk of absence	e^γ	Asymptotic t ratio	95% confidence interval		Log Likelihood
7 (± 3)	0.94	-0.57	0.76	1.16	-82317
14 (± 3)	1.12	1.13	0.92	1.37	-82317
21 (± 3)	1.28	2.37	1.04	1.56	-82314
28 (± 3)	1.44	3.51	1.18	1.77	-82311
35 (± 3)	1.27	2.27	1.03	1.56	-82315
42 (± 3)	1.10	0.85	0.88	1.37	-82317
49 (± 3)	0.91	-0.72	0.71	1.17	-82317

Note: Cox-Proportional estimates of the factor e^γ by which the hazard ratio of an absence for females relative to males increases in different sets of days after a previous absence episode (see equation 1). The analysis is restricted to females younger than 45. The row for 28 ± 3 corresponds to the second row and second column of Table 2.