

Biological Gender Differences, Absenteeism and the Earnings Gap

Andrea Ichino

UNIVERSITY OF BOLOGNA

Enrico Moretti *

UNIVERSITY OF CALIFORNIA, BERKELEY

May 20, 2008

Abstract

In most countries, women are more frequently absent from work than men. Using personnel data, we find that the absences of women below the age of 45 follow a 28-day cycle, while the absences of men and of women over 45 do not. We interpret this as evidence that the menstrual cycle increases female absenteeism. To investigate the effect on women's earnings, we use a simple model of statistical discrimination which suggests that absenteeism is more informative about shirking for men than for women. Consistent with the model, we find that absenteeism has a more negative effect on men's earnings and that this difference declines with seniority. The increased absenteeism induced by the 28-day cycle explains at least 14% of the earnings gender differential.

JEL-Code: J7, M5

Keywords: absenteeism, gender differentials.

In most countries, the earnings of female workers are lower than the earnings of male workers with similar observable levels of human capital and individual characteristics. In the US the conditional gender gap for white-collar workers is approximately -20 percent. In other European countries it is about -17 percent. A large literature has documented these earnings differences and analyzed several possible explanations.¹

One continuing limitation of this literature, however, is that it is unclear whether the factors that are proposed to explain gender differences in labor market outcomes are truly exogenous, or are instead endogenous responses to the fact that men and women may be treated differently in the labor market. For instance, starting with the seminal work by

¹For a recent survey of the literature, see Joseph Altonji and Rebecca Blank (1999).

Jacob Mincer and Solomon Polachek (1974), it has been well known that differences in labor market experience between men and women can account for a substantial share of the gender gap in earnings. This difference in labor market experience is attributed by some to biological differences between men and women. It is clear, for example, that it falls upon female workers to take time off work to give birth to children and to breast-feed them in the first few months of their lives. But an alternative explanation highlight the role of gender discrimination in the workplace and cultural biases in the allocation of family duties. For example, in the presence of returns to specialization in the labor market and home production, gender discrimination can result in substantial differences in labor market experience between men and women, even in the absence of any inherent gender difference. Empirically, it has proven difficult to distinguish between truly exogenous differences between men and women, and the endogenous responses to gender-specific labor market conditions.

A deeper assessment of the role of biology requires understanding whether there are truly exogenous biological differences between men and women that might explain some fraction of the male-female difference in earnings. In this paper, we focus on absenteeism. In most Western countries absenteeism is higher among female workers than among male workers. For example, in Europe, women take approximately 7.6 more sick-days per year than men with the same age, occupation and level of education. In the US and Canada, the corresponding figures are 3.1 and 5.2 days. While family-related commitments explain part of this gender gap in absenteeism, even among unmarried workers with no children women still take significantly more sick-days than men. Our findings suggest that part of this gender difference in absenteeism may be attributed to a biological difference between men and women and that this difference has small but non-trivial consequences for women careers and earnings. This is one of the first empirical studies to uncover a direct role of biological factors in the explanation of gender differences in labor market outcomes. We stress that our findings clearly do not rule out the importance of other factors that might be responsible for gender differences in outcomes, like gender discrimination or cultural biases.

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 most recent medical literature. The incidence of the 28-day cycle is no less pronounced for those workers up for promotion, who arguably have stronger incentives to 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.

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

The key insight of the model is that if male absenteeism depends only on the propensity to shirk and non-menstrual health shocks while female absenteeism is also driven by the menstrual cycle, then absenteeism is a noisier signal of worker quality for females than for males. If this is the case, signal extraction of underlying shirking rates based on absenteeism is 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, an absence episode is associated with a smaller earnings loss for women than men. Second, we find the same difference in slope

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

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.

Women in our sample earn about 13.5 percent less than men, conditional on their demographic characteristics. 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. The key identifying assumption for this counterfactual exercise is that the difference in unobserved ability between women and men does not decline with absenteeism. This assumption is consistent with the theoretical model, and is supported by the the empirical evidence on the predictions of the model.

We find that in the absence of 28-day cyclical absenteeism, the conditional gender gap in earnings would decline from -13.5 percent to -11.6 percent, a 14.1 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 15.3 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 that benefit women. Forcing employers, rather than women, to bear the monetary burden associated with menstruation may be counterproductive for the employment of women. But it is in theory possible to alleviate the cost of menstrual-related absenteeism using a gender-specific wage subsidy financed out of general taxation. A wage subsidy that favor female workers would shift part of the costs of menstrual-related absenteeism from women to men. The estimates presented in this paper

could in principle be used to quantify the magnitude of such subsidy. Because this is not a case of market failure, the rationale for the subsidy would be redistribution rather than efficiency. Whether society should address this biological difference with a gender-based wage subsidy depends on voters' tastes for redistribution. This conclusion is consistent with recent research that supports lower tax rates for women on fiscal efficiency grounds.³

The paper proceeds as follows. In Section I we test whether menstrual symptoms increase women's absenteeism. In Section II we test the predictions of a simple model of wage determination to investigate how the cost of an absence varies between men and women. In Section III 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 IV we conclude.

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

In many market economies, 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.⁴

In this Section we are interested in whether this gender difference in absenteeism may be caused by a specific biological factor: the menstrual cycle which affects women but not men. We test whether women's absences from work display a systematic 28-day cycle and we quantify what fraction of gender differences in absenteeism is due to absenteeism with cycle 28 days. We begin by showing some graphical evidence (subsection I.A) and then present more formal parametric tests (subsection I.B). In subsection I.C we calculate the number

³See, among others, Alberto Alesina, Andrea Ichino and Loukas Karabarbounis (2007).

⁴We are not the first to document that women have higher levels of absenteeism than men. See, for example, Lynn Paringer (1983), J.Paul Leigh (1983), Tim Barnby, Chriss D. Orme and John C. Treble (1991), Audrey VandenHeuvel and Mark Wooden (1995), Jessica P. Vistnes (1997), Sarah Bridges and Kate Mumford (2000). The literature has not provided convincing evidence on what the causes and consequences of these gender differences may be.

of absences due to the cycle 28. Finally, we tests whether the incidence of the 28-day cycle varies as a function of incentives in the workplace (subsection I.D).

I.A 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.⁵ We therefore exclude all employees who took maternity leave at any point during this period.⁶ This provides a sample of 16,208 workers. We focus on the 14,857 who have at least one illness-related absence during the three years observed. 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 in the managerial ranks.⁷

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,

⁵Under 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.

⁶We also exclude the 166 top managers, of whom only two are women.

⁷Given 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 Giovanni Maggi (2000), Ichino, Michele Polo and Enrico Rettore (2003) and Ichino and Regina Riphahn (2004).

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, Figure 2 repeats this exercise, now including all possible pairs of absences. Specifically, the figure shows the female-male difference in the distribution of number of days between the beginning of each absence calculated for all possible pairs of absences for workers with 2 or more absences. This figure illustrates that the probability that any two episodes are 28 days apart is higher for women than for men. The spike at 28 days is driven primarily by younger workers, and disappears with age. The top-right panel, which includes only workers under 45, displays a marked difference at 28 days. This difference disappears in the bottom-left panel, which includes workers 45 to 55, and in the bottom-right panel, which includes only workers 55 or older. This pattern is consistent with the timing of menopause.⁸

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. Three features of these figures warrant comment. First, the hazard is almost always higher for women, mirroring their higher overall absence rates. As discussed above, 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 percent and 35 percent of female and male absences begin on Monday, respectively. By comparison, the fraction of absences that

⁸The medical literature indicates that although many women experience menopause between 45 and 55, the age at onset varies greatly.

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

begin on other days of the week ranges from 11 to 21 percentage points for females, and 12 percent to 19 percent 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

I.B 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 David R. Cox (1972). Looking at two consecutive absences, we specify the hazard of the second as:

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

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; S_{it} is a dummy taking value of 1 if t is 7 or multiples; and Z is a vector of covariates.¹¹

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.

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 $\frac{h(t, X_{jt}, \Psi)}{\sum_{i=j}^N h(t, X_{it}, \Psi)} = \frac{e^{\alpha + \beta F_j + \gamma M_{jt} F_j + \delta S_{jt} F_j + \theta Z_j}}{\sum_{i=j}^N e^{\alpha + \beta F_i + \gamma M_{it} F_i + \delta S_{it} F_i + \theta Z_i}}$. 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.

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 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.

Table 2 presents estimates of the parameters e^β , e^γ and e^δ . The estimated coefficients are reported in the form of hazard ratios (with the t-statistics in parentheses). In the first panel, $e^\beta = 1.37$ indicates that the hazard of an absence from work is on average 35 percent higher for women younger than 45 than for men in the same age bracket. In addition to this higher overall risk, young women experience a higher incidence of absences with cycle 28 days. Specifically, the fact that $e^\gamma = 1.15$ for women under 45 and is statistically significantly different from 1 indicates that the hazard of an absence is 15 percent higher relative to males at cycles of 28 days. The estimate of e^δ suggests that the hazard of an absence occurring 7 days after a previous episode is 6 percent lower for women than for men and that this difference is statistically significant. Thus, young women are significantly more likely than young men to be absent every 28 days, while the opposite is true for absences 7 days apart.

For workers older than 45, the overall gender difference in column 1 increases slightly, indicating that the higher absenteeism of women relative to men does not disappear with age. By contrast, we see no difference across genders in the incidence of the 28-day pattern in column 2. The coefficient is statistically indistinguishable from one.

To check that these results are not driven by the 7-day periodicity, we have also estimated the same models restricting the sample to those spells that do not begin on a Monday. These estimates are reported in the second panel of Table 2. The point estimate of the female-male difference in the incidence of the 28-day cycle for young workers is virtually unchanged, although the statistical significance declines slightly due to the smaller sample. The estimate of e^δ is now close to 1 for both young and old workers, indicating that the gender difference in the 7-day cycle for young workers is mostly due to Monday absences.

In columns 4, 5 and 6, we repeat the analysis 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 is 15-17 percent higher relative to males at cycles of 28 days. As before, we find no such effect for those over 45.

The estimates in Table 2 are obtained by focusing on a cycle of 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 different than 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 28-day cycle – which corresponds to the estimate in the first row and second column of Table 2 – is the largest and the most precise. The other coefficients and corresponding log likelihoods decline almost monotonically as we move further away from 28. These results confirm the visual impression obtained from Figures 1 and 2.¹²

I.C 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, in view of the fact that menstrual symptoms are unlikely to induce

¹²We tried to obtain data on pill use to see whether the incidence of the 28-day 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-day 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.

long absences. Moreover, we call two absences cyclical if they are both short, and they are between 26 and 30 days apart, or multiples thereof. This because menstrual cycles vary enormously, both across women and across months for a given woman, and Figures 1 and 2 are consistent with this notion.¹³

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.¹⁴

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.¹⁵

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

¹³Mitchell D. Creinin, Sharon Keverline and Leslie A. Meyn (2004) report that for 46 percent of their subjects the length of the cycle can vary by 7 days or more. For 20 percent the length of the cycle can vary by 14 days or more.

¹⁴Our 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.

¹⁵Furthermore, in some of our models, we use a more conservative definition of cyclical absences that effectively puts an upper bound on the number of false positives. Specifically, in some models we re-classify a cyclical absence as non-cyclical if we find another absence exactly 7 days before or after that day. Our estimates are not sensitive to this re-classification.

cyclical absences. It is apparent that the distribution for women stochastically dominates the distribution for men.¹⁶

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 percent 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 percent 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. Two features are interesting. First, the distribution for women is clearly shifted to the right of that for men. For example, the fraction of men and women for whom the number of estimated cyclical absences is 0 is 55 percent and 29 percent, respectively. Second, the heavy menstrual symptoms are concentrated among a relative few women. The vast majority of women has 0 or few cyclical absences. For example, more than two third of women have only 0, 1 or 2 cyclical absences

¹⁶One 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 percent of all pairs of absences is a false positive for men, while the corresponding number for women is 3.4 percent). We conclude that this problem is unlikely to affect our estimates in any significant way.

per year. Furthermore, of these absences some are likely to be "false positives" due to the 7-day cycle, as indicated by the fact that some men also have cyclical absences. By contrast, 5 percent of women have 9 or more days of cyclical absence per year (almost one day per month). Among the men, only 1 percent has 9 or more days of cyclical absence per year.

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 percent of premenopausal women regularly experience some form of mild premenstrual symptoms.¹⁷ A smaller fraction of women typically meet all the criteria for the clinical definition of premenstrual syndrome, or for its more severe version, 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.¹⁸ Our estimates in Table 3, based on the same age range, indicate that the average

¹⁷See, among others, Susan R. Johnson (1987), Patricia A. Deuster, Tilahun Adera and Jeannette South-Paul (1999), Anita Chawla et al. (2002) and Barbara Sternfeld 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, Adera and South-Paul 1999, p. 122).

¹⁸They report that 17.3 percent of their sample had "severe" symptoms. For 5 percent 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 percent 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: $(1.1 * 0.17 + 1.3 * 0.05) * (365 / (28 * 2)) = 1.7$.

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.

I.D Incentives and Work Environment

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.

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.

To address this issue, in Table 5 we focus on workers who received a promotion during the period of observation, 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.¹⁹ 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.²⁰ This is not surprising, since workers have strong incentives to minimize their absenteeism in the months leading

¹⁹We 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.

²⁰The 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.

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 5 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.

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 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, Table 5 is consistent with the notion that the periodicity in absenteeism is not lower when the cost of such absenteeism is high. This evidence seems 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. 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.

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 this questions, we have modified model 1 to allow for a differential incidence of the 28-day cycle depending on the fraction of women employed in the branch or on whether the relevant branch is in the South or in the North.

We find that a prevalently female work environment is associated with substantially higher overall absenteeism, but a prevalently female work environment is not associated with an increase in the effect of menstrual cycles. Similarly, we find that Southern branches

are characterized by significantly higher absenteeism than branches in other regions, but 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. See Ichino and Enrico Moretti (2006) for the estimates.

Overall, this indicates that the association between the menstrual cycle and absenteeism does not change between environments with high or low rates of overall absenteeism.

II The Price of Absenteeism

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.

In our sample of Italian bank workers, women earn less than the men. The average yearly earnings for women and men are 25,020 and 29,034 Euros, respectively. 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 -20 percent. In other European countries it is about -17 percent.

How much of the observed gender gap in earnings is explained by the additional absenteeism generated by the menstrual cycle? The answer to this question depends on the cost for workers in terms of reduced earnings of an additional day of absence. In this firm (and in most Italian firms) workers receive a fixed monthly salary, irrespective of the number of days of absence in that specific month.²¹ Although in the short run we do not expect the salary of a worker to adjust to the number of absences in a given month, in the long run we expect employers to reward workers who tend to have low absenteeism. This may occur through direct merit-based wage increases, or, most likely, through faster promotion of employees with lower absenteeism. While union rules constrain the ability of differentiating compensation within a defined occupational level, Italian employers have flexibility in deciding on merit-based promotions and job assignments.

²¹For short absences (3 days or less), the employer is fully responsible for paying for the workers salary. For longer absences (more than 3 days), workers wages are paid in part by Social Security.

The *direct* effect of menstrual-related absenteeism can be calculated by estimating the value of work time lost due to the menstrual cycle:

$$\frac{\text{Days of work lost due to cycle} \times \text{Women's average daily earnings}}{\text{Gender gap in earnings}} \quad (2)$$

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 the average earnings of women and men, 4.4 percent of the earnings gap can be explained by the direct effect of this absenteeism on earnings: $[1.5 \times (25020/214)] / (29,034 - 25,020) = 4.4\%$.

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, 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. Second, in most white-collar jobs the cost to the employer of the disruption caused by an unplanned day of absence is surely more than the daily earnings of that person.²²

Third, and most importantly, the calculation above does not reflect the signaling value of avoiding absences. When worker productivity is not perfectly observed, 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. There are reasons to expect that this signaling cost differ significantly for men and women.

In the rest of this section, we describe a simple model that clarifies how the relationship between absenteeism, worker quality, and earnings differs for men and women. Our model provides a set of testable implications that we bring to the data. The predictions of the model

²²Moreover, 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. For example, in a clinical study, Chawla et al. (2002) estimate that women with severe PMS symptoms experience decreases in productivity of 48.2 percent — 64.4 percentage points 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 (Jean Endicott and John Nee 1997) and time diaries. All the differences are statistically significant at the 1 percent level. Self-assessed productivity declines were between 13.8 and 22.7 percentage points. Yan Chen, Peter Katuscak and Emre Ozdenoren (2005) find evidence of lower performance for women during menstruations in a laboratory experiment.

become useful in the next Section, when we use variation across workers in the incidence of the 28-day cycle to quantify the total 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.

II.A Gender Differences in the Relationship Between Absenteeism and Earnings: Theory

A formal model is developed in the Appendix. Here we provide the intuition and the testable implications. If employers cannot directly observe individual productivity, they might use observable worker characteristics, including absenteeism, to predict productivity and set wages.²³ We assume that employers set wages according to a simple model of statistical discrimination, weighting their gender-specific priors and the observed signal (absenteeism). In particular, the employer's best guess of the unobserved propensity to shirk of a worker is a weighted average of prior and the signal, with weights that reflect the relative precision of prior and signal.

The key insight is that if menstrual-related absences are not a signal of shirking, absenteeism is a noisier signal of shirking attitudes for females than for males and therefore observed absenteeism has a larger effect on employer's priors for men. In other words, signal extraction based on absenteeism is more informative about shirking for males than for females. This is due to the fact that for men, a day of absence reflects either a non-menstrual related health condition or shirking; while for women, a day of absence reflects one of those factors or menstrual cycle.

The key prediction of the model is that while we should expect the relationship between earnings and absenteeism to be negative for both males and females (as more absenteeism implies more shirking), this relationship should be more negative for males than for females. This implies that an absence episode is associated with a smaller earnings loss for women than men. A second important prediction is that this gender difference in slope should decline with seniority, since the informational content of absenteeism declines as employers

²³See for example Bengt Holmstrom (1999) and Dennis J. Aigner and Glenn G. Cain (1977). For an empirical example of the link between absenteeism and shirking see Peter Skogman Thoursie (2004). In a different context, Paul Milgrom and Sharon Oster (1987) use asymmetric information to explain gender wage differences.

learn more about their workers.²⁴ On the other hand, if menstrual-related absences were a signal of shirking, there should be no gender difference in the relationship between earnings and absenteeism at low or high levels of seniority.

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 model 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 slope in the relationship between earnings and absenteeism remains negative for both genders and smaller in absolute value for women. The main effect of introducing effort is that the gender gap in absenteeism widens, since women's equilibrium effort is lower.

A first testable implication on the relationship between earnings and cyclical absenteeism follows from the fact that absenteeism is a noisier measure of shirking for men than for women.

Proposition 1. 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.

Of course, this prediction applies only if menstrual episodes do not reflect shirking. As mentioned above, if menstrual episodes do reflect shirking, then in a regression of earnings on cyclical absenteeism the coefficient for females should be the same as the coefficient for males.²⁵

²⁴This last point has been made in a different context by Henry Farber and Robert Gibbons (1996) and Altonji and Charles Pierret (2001).

²⁵While the model does not distinguish between total absenteeism and cyclical absenteeism, in this paper we are interested in the latter. 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

A second prediction focuses on how the relationship between earnings and absenteeism varies over time. Over time, the true productivity of a worker gets revealed to the employer. This implies that gender-based statistical discrimination become less important over time. Therefore, any gender difference in the relation between earnings and absenteeism should disappear with seniority.

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, if menstrual episodes do not reflect shirking, 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.²⁶

II.B Gender Differences in the Relationship Between Absenteeism and Earnings: Evidence

We now take the three predictions to the data. We note that the predictions from the model do not necessarily involve causality, since they are simply equilibrium outcomes.

(1) Evidence on Earnings and Careers. The entry in the first column in the top panel of Table 6 shows that the unconditional earnings gap in our sample is -20 percent.²⁷

When we control for a quadratic in age and the number of non-cyclical absences in column

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-day cycle.

²⁶The first part of Proposition 3 is easily derived using equation 9 in the Appendix. The second part of Proposition 3 derives from the fact that under our assumptions, absenteeism is a stationary variable and its correlation with the time invariant propensity to shirk should not change over time.

²⁷In this and all the remaining Tables, we use only workers who are 45 or younger.

2, the earnings gap declines to -13.5 percent. 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. 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 percent. The cost for female workers is only 1.5 percent. 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.

We have re-estimated the model in column 3 using a more conservative definition of cyclical absences. In particular, we have re-classified a cyclical absence as non-cyclical if we find another absence exactly 7 days before or after that day. Our estimates are not sensitive to this re-classification.²⁸ As a second specification check, in column 4 we add the interaction of the female dummy and non-cyclical absences. If non-cyclical absenteeism has the same correlation with propensity to shirk for males and females, the coefficient on the interaction of female and non-cyclical absences should be zero. On the other hand, it is possible that for an employer, an additional day of non-cyclical absenteeism is less likely to be a signal of shirking for females than males. This may occur, for example, if it falls more upon female workers than male workers to take days off to deal with family commitments, such as when there are sick children or sick parents to care for, or when there are parent-teacher conferences, or for other types of family duties. In these cases, non-cyclical absenteeism has a lower correlation with propensity to shirk for females than males, and the coefficient on the interaction of female and non-cyclical absences should be larger than zero, although smaller than the coefficient on the interaction of female and non-cyclical absences. Our estimate in column 4 indicates that the coefficient on the interaction is positive, although significantly smaller than the coefficient on the interaction of the female dummy and cyclical absences.

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

²⁸For example, the coefficients on cyclical absences and on cyclical absences \times female are, respectively, -0.024 (0.001) and -0.011 (0.002), respectively. This reclassification effectively puts an upper bound on the number of false positive described in Section I.C.

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 percent less likely to be promoted to a management position (column 1), or 11 percent 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. The mean (std deviation) of the dependent variable is 6.1 (2.2). Again, the estimates shown in Table 6 are consistent with Proposition 1.²⁹

(2) Evidence on Employer Learning. 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, attrition is negligible.³⁰ The specification in Table 7 generalizes the one in column 3 of Table 6 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

²⁹Alternative 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 firms tasks are fairly homogenous. For instance, the tasks performed by clerks in most branches of this bank are quite standardized.

³⁰Only 12 workers quit and 4 were fired during the 3-year period for which we have data. An additional 4 workers died and 261 retired. This low attrition is due to good working conditions and high salaries.

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 6, 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 -0.032 for men and only -0.019 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: -0.007 (0.0003). 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.³¹

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.

(3) Evidence on 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 8 we use schooling. The results are consistent with Proposition 3. For men, increases in cyclical absenteeism are associated with a steep decline in schooling.

³¹Note that the effect of seniority on earnings is not immediately apparent, because in the regression we 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.

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.

Similarly, in column 2 the dependent variable is non-cyclical absenteeism. We find that consistent with Proposition 3, workers with more cyclical absences are more likely to have non-cyclical absences, but this is significantly less true for women than for men. In column 3 the dependent variable is a dummy equal to 1 if the worker is involved in any misconduct episodes in the three years observed.³² 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 4 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 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 5), 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 7, using as dependent variables the four indicators of workers quality that we use in Table 8. In all cases, we found that the coefficient on the triple interaction between cyclical absence, seniority and gender is statistically insignificant.³³

³²These 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.

³³The coefficient (std. error) on the triple interaction for schooling is 0.0005 (0.003); for misconduct,

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

In this Section, we seek to estimate the counterfactual gender gap in earnings in the absence of menstrual absenteeism. In the previous Section, we have provided an estimate of the direct effect of the menstrual-related absenteeism on the gender gap in earnings. This estimate is simply the fraction of the gender gap in earnings that is due to work time lost due to menstrual cycle (Equation 2). We have argued that such estimate is likely to capture only part of the cost of absenteeism, and therefore should be considered a lower bound.

We now adopt an alternative approach to determine the total effect of menstrual cycle on the gender gap in earnings. To do so, we need to estimate the counterfactual earnings level of women in the absence of menstrual absenteeism. We construct a counterfactual earnings level of women in the absence of menstruation by assigning the male distribution of absenteeism to females. Specifically, 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. The counterfactual gender gap that emerges from this procedure can also be interpreted as the term of an Oaxaca decomposition (Ronald Oaxaca 1973) that originates from fixing the cost of cyclical absenteeism across genders and letting wages differ between males and females only because of their different propensity to be cyclical in absenteeism. Notably, the key identifying assumption for this counterfactual exercise follows from the theoretical model described in Section II.

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. 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:

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

0.0004 (0.0004); for vacation, 0.0001 (0.004); for strike, -0.0009 (0.0009).

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:

$$\tilde{Y}_f - \tilde{Y}_m = \pi_{1m}(Y_{1f} - Y_{1m}) + \pi_{2m}(Y_{2f} - Y_{2m}) + \pi_{3m}(Y_{3f} - Y_{3m}) \quad (4)$$

The counterfactual earnings gap is simply a weighted average of the wage gap in three groups, with weights reflecting the male distribution across groups. This approach relaxes the linearity assumption implicit in the models in Table 6, allowing the relationship between earnings and absenteeism to be non-linear. Intuitively, equation 4 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.

This strategy provides the valid counterfactual gender gap under two assumptions. First, 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. Second, the female-male difference in unobservables is the same for all three groups, or at least it does not decline with cyclical absences:

$$\begin{aligned} E(e_i|F, j = 3) - E(e_i|M, j = 3) &\geq E(e_i|F, j = 2) - E(e_i|M, j = 2) \\ &\geq E(e_i|F, j = 1) - E(e_i|M, j = 1) \end{aligned} \quad (5)$$

where e_i is worker i unobserved ability. 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.

This assumption is plausible because it follows from the simple model in Section II. Specifically, 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. Notably, the empirical validity of the model and therefore of the identifying assumption is supported by the evidence in Tables 6 to 8.

Equation 4 requires us to estimate the gender difference in earnings for the three groups. We do so by fitting the following equation:

$$\log Y_i = \beta_1 + \beta_2 C_{2i} + \beta_3 C_{3i} + \gamma_1 C_{1i} F_i + \gamma_2 C_{2i} F_i + \gamma_3 C_{3i} F_i + \mu X_i + e_i \quad (6)$$

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 $\gamma_3 = (Y_{3f} - Y_{3m})$.

Column 1 in Table 9 shows estimates of the β 's and γ 's. Entries in rows 1 and 2 indicate that male earnings decline as we move from group 1 to groups 2 and 3. This is not surprising, given that we already know that earnings decline with cyclical absenteeism. Entries in rows 3, 4 and 5 show that—consistent with Proposition 1 and with Table 6—the decline is steeper for men than for women. This can be seen from the fact that the coefficient γ_1 is more negative than γ_2 ; and γ_2 is more negative than γ_3 .

The bottom of the Table compares the observed gender gap with the counterfactual gender gap. The observed gender gap in earnings, conditional on covariates, is -13.5 percent. This number is reported for convenience from Table 6, column 2. To obtain the counterfactual gender gap, we re-weight our estimates of the γ 's, using the male distribution across the three groups as weights (equation 4). Specifically, we estimate that the counterfactual gender gap is -11.9 percent. 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.

The estimates in Table 9 are simply a more general version of the linear models in Table 6. The main advantage is that these estimates allow for a non linear relationship between earnings and cyclical absenteeism. We repeated the same exercise using an even finer grid. Specifically, we divided workers in 10 groups (instead of 3) based on their number of cyclical absences. This specification has the advantage to account more precisely for non-linearities. In this case, the estimated counterfactual gender gap is even lower: -11.6 percent. This estimate implies that if women did not experience 28-day cyclical absenteeism,

the earnings difference between females and males would be 14.1 percent lower than the observed difference.³⁴

Note that 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 “partial equilibrium” 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. Put differently, our counterfactual exercise answers the following question: if there were a medication capable to eliminate menstrual pain, what would happen if some women took this medication, but most other women did not? Interestingly, this scenario is not too far fetched. The Food and Drug Administration has just approved the first pill that is designed to eliminate menstrual periods (New York Times, April 20, 2007).

We can compare this estimate, 14.1 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 14.1 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.

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

³⁴Estimates based on linear models are slightly smaller, although not very different. The 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 equation 2. The only difference is that we have substituted “daily earnings” with “cost of a day of cyclical absence”. Our estimates in Table 6 suggest that a day of cyclical absence costs women 1.5 percent of earnings. Given that women earn on average 25,020 Euros, the formula implies that about 9.3 percent of the earnings gap can be explained by the direct effect of this absenteeism on earnings: $[1.5 \times (25,020 \times 0.015)] / 4014 = 9.3\%$. We thank Claudia Goldin for suggesting this calculation.

workers in 10 groups—instead of three— to better account for non-linearities, we find even larger effects: 15.3 percent for the probability of being promoted to manager and 33.3 percent for occupational level.

IV Conclusions

In most countries women take more sick days than men. We show 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 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.

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 percent. 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 -11.6 percent. In other words, the gender gap in earnings would be 14.1 percent lower. A similar calculation shows that the gender gap in the probability of being promoted to manager would be 15.3 percent 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.

Our findings may have significant policy implications. If one wanted to redistribute the cost of menstrual-related absenteeism from women to men, it would in principle be possible to adopt a gender-specific wage subsidy. A wage subsidy that targets women and is financed out of general taxation would shift part of the economic costs of menstruation from women to men. The estimates presented in this paper could in principle be used to set the magnitude of such subsidy. Of course, this is not a case of market failure and the rationale for the subsidy would only be redistributive. The rationale of such subsidy would therefore depend on voters' tastes for redistribution.

* We would like to thank Sascha Becker, Rebecca Blank, Ernesto Dal Bo, Janet Currie, Claudia Goldin, Stefano Gagliarducci, Bryan Graham, Soren Johansen, David Lee, Thomas Lemieux, 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, two anonymous referees and seminar participants at Berkeley, Bologna, CES-ifo, EUI, Houston, IFAU, IHS Vienna, Jerusalem, Linz, Milano, NBER Personnel Meeting, NBER Labor Meetings, Paris, Public Policy Institute of California, Rice, Siena, SOLE Meetings, Texas A&M, Torino and Warwick 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, Francesco Manaresi and Ashley Langer provided excellent research assistance. Email: andrea.ichino@unibo.it and moretti@econ.berkeley.edu.

References

- Alesina, Alberto, Andrea Ichino, and Loukas Karabarbounis.** 2007. "Gender Based Taxation and the Division of Family Chores." National Bureau of Economic Research Working Paper 13638.
- Aigner, Dennis J., and Glenn G. Cain.** 1977. "Statistical Theories of Discrimination in labor markets." *Industrial and labor Relations*, 30(2): 175–187.
- Altonji, Joseph, and Rebecca Blank.** 1999. "Gender and Race in the Labor Market." In *Handbook of Labor Economics*. ed. Orley C. Ashenfelter and David Card, Volume 3C: 3143–3259. New York, NY: Elsevier Science Press.
- Altonji, Joseph, and Charles Pierret.** 2001. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics*, 116(1): 313–350.
- Barmby, Tim, Chris D. Orme, and John C. Treble.** 1991. "Worker Absenteeism: An Analysis Using Microdata." *Economic Journal*, 101: 214-229.
- Bridges, Sarah, and Kate Mumford.** 2001. "Absenteeism in the UK: A Comparison Across Gender." *Manchester School*, 69(3): 276-284.
- Card, David, and Brian 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, Yan, Peter Katuscak, and Emre Ozdenoren.** 2005. "Why Can't a Woman Bid More Like a Man." CERGE-EI Working Paper 275.
- Creinin, Mitchell D., Sharon Keverline, and Leslie A. Meyn.** 2004. "How regular is regular? An analysis of menstrual cycle regularity." *Contraception*, 70(4): 289–292.
- Cox, David R.** 1972. "Regression models and life tables (with discussion)" *Journal of the Royal Statistical Society - series B*, 34: 187-220.
- Deuster, Patricia A., Tilahun Adera and Jeannette South-Paul.** 1999. "Biological, Social and Behavioural Factors Associated with Premenstrual Syndrome." *Archives of Family Medicine*, 8: 122–128.
- Endicott, Jean, and John Nee.** 1997. "Assessment measures for clinical studies. Endicott Work Productivity Scale: A new measure to assess treatment effects." *Psychopharmacological Bulletin*, 33: 3–6.
- Farber, Henry and Robert Gibbons.** 1996. "Learning and Wage Dynamics." *Quarterly Journal of Economics* 111(3):1007–47.
- Holmstrom, Bengt.** 1999. "Managerial Incentive Problems - A Dynamic Perspective." *Review of Economic Studies*, 66(1): 169–182.

- Ichino, Andrea, and Giovanni 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, Andrea and Enrico Moretti.** 2006. "Biological Gender Differences, Absenteeism and the Earning Gap." National Bureau of Economic Research Working Paper 12369.
- Ichino, Andrea, Michele Polo, and Enrico Rettore.** 2003. "Are Judges Biased by Labor Market Conditions?" *European Economic Review*, 47(5): 913–44.
- Ichino, Andrea, and Regina 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, Susan R.** 1987. "The Epidemiology and Social Impact of Premenstrual Symptoms." *Clinical Obstetrics and Gynecology*, 30:367–376.
- Leigh, J. Paul.** 1983. "Sex Differences in Absenteeism." *Industrial Relations*, 22(3): 349-361.
- Milgrom, Paul, and Sharon Oster.** 1987. "Job Discrimination, Market Forces, and the Invisibility Hypothesis." *Quarterly Journal of Economics*, 102(3): 453-476.
- Mincer, Jacob, and Solomon Polachek.** 1974. "Family Investment in Human Capital: Earnings of Women." *Journal of Political Economy*, 82(2): S76-108.
- Oaxaca, Ronald.** 1973. "Male-Female Wage Differentials in Urban Labor Markets." *International Economic Review*, 14: 673-709.
- Paringer, Lynn.** 1983. "Women and Absenteeism: Health or Economics?" *American Economic Review*, 73(2): 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." *Obstetrics and Gynecology*, 99: 1014-1024.
- VandenHeuvel, Audrey, and Mark Wooden.** 1995. "Do Explanations of Absenteeism Differ for Man and Women?" *Human Relations* 48(11): 1309-1329.
- Vistnes, Jessica P.** 1997. "Gender Differences in Days Lost from Work due to Illness." *Industrial and Labor Relations Review* 50(2): 304–322.

Appendix: A Simple Model of Statistical Discrimination

(A) Exogenous Effort Assume that the productivity *per unit of working time* of employee i is given by

$$Y_i = c - S_i \quad (7)$$

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

$$W_{it} = E(Y_i|X_{it}) \quad (8)$$

Because 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.³⁵ 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. In particular, we assume that

$$X_{it} = S_i + \mu H_{it} \quad (9)$$

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. If this were not the case μ would be the same for both genders.³⁶

We also assume that

$$S_i \sim N\left(\omega, \frac{1}{p}\right) \quad (10)$$

$$H_{it} \sim N\left(\eta, \frac{1}{q}\right) \quad (11)$$

where the parameters ω, η, p, q and μ are known to everyone.

In period 1, employers use the Normal Learning Model to predict which workers are productive and which workers are shirkers, based on the level of observed absenteeism:

$$\begin{aligned} E(S_i|X_{i1}) &= E(S_i + \mu\eta|X_{i1}) - \mu\eta \\ &= \frac{p}{p + \frac{q}{\mu^2}}\left(\omega - \frac{q\eta}{\mu p}\right) + \frac{\frac{q}{\mu^2}}{p + \frac{q}{\mu^2}}X_{i1} \end{aligned} \quad (12)$$

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 ($\omega - \frac{q\eta}{\mu p}$). Unlike the most commonly used models of statistical discrimination, in our model the crucial difference

³⁵In reality, absenteeism also induces a mechanical loss of output. For notational simplicity, we focus on signalling and ignore the mechanical loss of output. Including this mechanical effect does not change our results.

³⁶Note 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 a cycle of 28 days.

between genders is not the difference in the mean of the prior, but in the variance of the prior. The wage paid by the firm in period 1 is

$$W_{i1} = c - E(S_i|X_{i1}) = \alpha - \beta X_{i1} \quad (13)$$

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

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

It is easy to see that

$$\beta_{male} > \beta_{female} \quad (15)$$

We now determine how relationship between earnings and absenteeism evolves over time, as employers learn more about each worker's quality, S_i . Iterating the Normal Learning Equation, we can see that after t periods Equation 12 becomes:

$$E(S_i|X_{i1}, \dots, X_{it}) = \frac{p}{p + t\frac{q}{\mu^2}}(\omega - \frac{q}{\mu}\eta) + \frac{\frac{q}{\mu^2}}{p + t\frac{q}{\mu^2}} \sum_{s=1}^t X_{is} \quad (16)$$

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

$$W_{it} = \alpha_t - \beta_t \bar{X}_{it} \quad (17)$$

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

$$\beta_t = \frac{\frac{q}{\mu^2}}{\frac{p}{t} + \frac{q}{\mu^2}} \quad (18)$$

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. (The slope does not go to zero with perfect information, because workers with high absenteeism have, by assumption, a higher propensity to shirk.)³⁷

³⁷If one could measure S_i , an hypothetical regression of S_{it} on X_{it} would yield a slope coefficient equal to

$$\frac{\text{cov}(S_i, X_{it})}{\text{var}(X_{it})} = \frac{\text{var}(S_i)}{\text{var}(S_i) + \mu^2 \text{var}(H_{it})} = \frac{\frac{1}{p}}{\frac{1}{p} + \frac{\mu^2}{q}} \quad (19)$$

because S_i and H_{it} are orthogonal. Since μ is larger for females than males, equation 19 implies a steeper positive slope for males if the dependent variable is the propensity to shirk S_i . Note that in this case the OLS coefficient in equation 19 is identical to the parameter β of the Normal Learning Model in equation 14. 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.

(B) Endogenous Effort This 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. 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 9 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 .³⁸ 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

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

Compared with equation 13, the slope coefficient β is unchanged. In particular, it remains steeper for men than women. 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 13. In particular, the difference between the intercept for men and women is now larger because women have a lower incentive to exert effort.³⁹

³⁸For simplicity, we assume that θ_i is orthogonal to S_i and H_{i1} .

³⁹These findings do 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.

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: Standard errors in parentheses. 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. 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 risk of a menstrual cycle.

	Without controls			With controls		
	e^β	e^γ	e^δ	e^β	e^γ	e^δ
	(1)	(2)	(3)	(4)	(5)	(6)
<u>By Age Group</u>						
Under 45	1.37	1.15	0.94	1.39	1.15	0.95
	(36.14)	(2.17)	(-2.22)	(35.94)	(2.16)	(-2.04)
Above 45	1.42	0.99	0.97	1.29	0.99	0.97
	(19.71)	(-0.08)	(-0.55)	(13.82)	(-0.08)	(-0.43)
<u>Excluding Mondays</u>						
Under 45	1.34	1.18	0.97	1.36	1.17	0.98
	(28.28)	(1.96)	(-0.81)	(28.47)	(1.95)	(-0.69)
Above 45	1.39	0.98	1.01	1.26	0.98	1.01
	(15.81)	(-0.09)	(0.20)	(10.95)	(-0.09)	(0.28)

Notes: Asymptotic t-ratios in parentheses. Entries are the Cox-Proportional Hazard ratios for the occurrence of an absence episode with time measured from the beginning of the previous episode, for all episodes observed between January 1, 1993 and December 31, 1995. The total number of spells is 97637. Results 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; e^δ is the factor by which the hazard ratio of females relative to males increases every seven days. An hazard ratio equal to 1 indicates absence of effect. A ratio larger (smaller) than 1 indicates a positive (negative) effect. In columns 4 to 6, results are conditional on age, years of schooling, marital status, number of children, managerial occupation, seniority and dummies for the weekday in which the spell begins.

Table 3: Absenteeism, by Type and Gender

	Men	Women	Difference		
	(1)	(2)	Unconditional (3)	Conditional (4)	Conditional (5)
<u>Days of Illness-Related Absence</u>					
Total Number of Days in a Year	8.2	12.9	4.6 (0.3)	5.2 (0.3)	5.4 (0.3)
Estimated Number of Cyclical Days in a Year	1.3	2.9	1.4 (0.06)	1.5 (0.06)	1.5 (0.06)
<u>Episodes of Illness-Related Absence</u>					
Total Number of Episodes in a Year	2.1	3.6	1.5 (0.5)	1.6 (0.5)	1.6 (0.6)
Estimated Number of Cyclical Episodes in a Year	0.9	2.0	1.1 (0.04)	1.1 (0.04)	1.1 (0.04)
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	56	29
1	20	23
2	10	15
3	5	9
4	3	6
5	2	5
6	1	4
7	1	2
8	1	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, before and after a promotion

	Before promotion			After promotion		
	e^β	e^γ	e^δ	e^β	e^γ	e^δ
<u>By Age Group</u>						
Under 45	1.52 (3.48)	3.87 (2.91)	0.60 (-0.98)	1.75 (4.34)	2.87 (2.03)	0.93 (-0.15)
Above 45	1.70 (1.55)	0.00 (0.00)	6.18 (2.25)	1.87 (1.66)	0.00 (0.00)	3.60 (1.60)
<u>By Age Group, with Controls</u>						
Under 45	1.46 (2.81)	3.77 (2.85)	0.59 (-1.03)	1.64 (3.42)	2.89 (2.04)	0.94 (-0.12)
Above 45	1.54 (1.17)	0.00 (0.00)	6.25 (2.25)	1.94 (1.60)	0.00 (0.00)	2.32 (0.95)

Notes: Asymptotic t-ratios in parentheses. The sample includes only workers who received a merit promotion between 1993 and 1994. Entries are Cox-Proportional Hazard ratios for the occurrence of an absence episode with time measured from the beginning of the previous absence episode. Note that an hazard ration equal to 1 indicates absence of effects. 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; e^δ is the factor by which the hazard ratio of females relative to males increases every seven days. 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. 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 6: Earnings and Career Equations - Linear Models

	(1)	(2)	(3)	(4)
<u>Model 1: Earnings</u>				
Female	-0.204 (0.006)	-0.135 (0.006)	-0.144 (0.008)	-0.153 (0.008)
Cyclical Absences			-0.025 (0.001)	-0.023 (0.001)
Female \times cyclical absences			0.010 (0.002)	0.008 (0.002)
Female \times non-cyclical absences				0.002 (0.0007)
<u>Model 2: Promoted to Manager</u>				
Female	-0.183 (0.009)	-0.111 (0.009)	-0.138 (0.012)	-0.151 (0.013)
Cyclical Absences			-0.029 (0.002)	-0.027 (0.002)
Female \times cyclical absences			0.017 (0.003)	0.015 (0.003)
Female \times non-cyclical absences				0.003 (0.001)
<u>Model 3: 13 Occupation Levels</u>				
Female	-0.754 (0.050)	-0.216 (0.047)	-0.368 (0.060)	-0.473 (0.067)
Cyclical Absences			-0.184 (0.010)	-0.168 (0.010)
Female \times cyclical absences			0.108 (0.015)	0.088 (0.017)
Female \times non-cyclical absences				0.022 (0.005)
Controls for non-cyclical absences	N	Y	Y	Y
Controls for age	N	Y	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 7: The Relationship Between Earnings and Cyclical Absences, by Gender and Firm Seniority

	Dependent Variable is Earnings (1)
Female \times Cyclical Absences \times Seniority	-0.0007 (0.0003)
Female \times Cyclical Absences	0.013 (0.002)
Female \times Seniority	0.002 (0.001)
Cyclical Absences \times Seniority	0.001 (0.0002)
Female	-0.145 (0.011)
Cyclical Absences	-0.032 (0.001)
Seniority	-0.011 (0.000)

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

	Schooling	Non-Cyclical Absenteeism	Misconduct	Days of Vacation Taken	Days of Strike
	(1)	(2)	(3)	(4)	(5)
Average Dep. Variable	13.1	7.53	0.09	19.5	0.97
Days of Cyclical Absence	-0.104 (0.014)	1.530 (0.057)	0.028 (0.001)	0.018 (0.015)	0.002 (0.003)
Days of Cyclical Abs. \times Female	0.094 (0.023)	-0.359 (0.098)	-0.018 (0.002)	-0.029 (0.025)	-0.000 (0.005)
Controls	Y	Y	Y	Y	

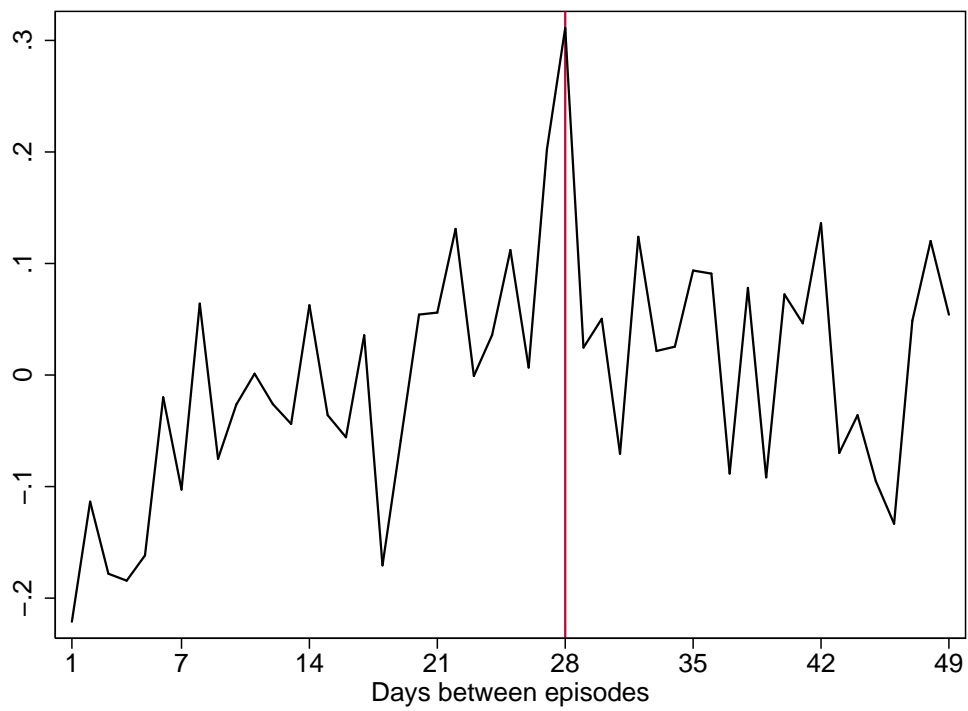
Notes: Standard errors in parenthesis. Controls include a dummy for females, a quadratic in age and dummies for number of days of non-cyclical absences. Control in column 2 do not include dummies for number of days of non-cyclical absences. Sample includes workers 45 or younger.

Table 9: 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)	-0.042 (0.007)	-0.037 (0.010)	-0.251 (0.053)
High number of cyclical absences (β_3)	-0.118 (0.007)	-0.135 (0.010)	-0.821 (0.053)
Small number of cyclical absences \times female (γ_1)	-0.131 (0.012)	-0.107 (0.019)	- 0.302 (0.092)
Medium number of cyclical absences \times female (γ_2)	-0.118 (0.012)	- 0.116 (0.018)	-0.222 (0.091)
High number of cyclical absences \times female (γ_3)	-0.099 (0.009)	- 0.059 (0.014)	0.142 (0.070)
Observed Gender Gap (conditional)	-0.135	-0.111	-0.216
Counterfactual Gender Gap (conditional)	-0.119	-0.096	-0.161
Percent of the Observed Gap "Explained" by Cycle	11.8%	13.5%	25.4%

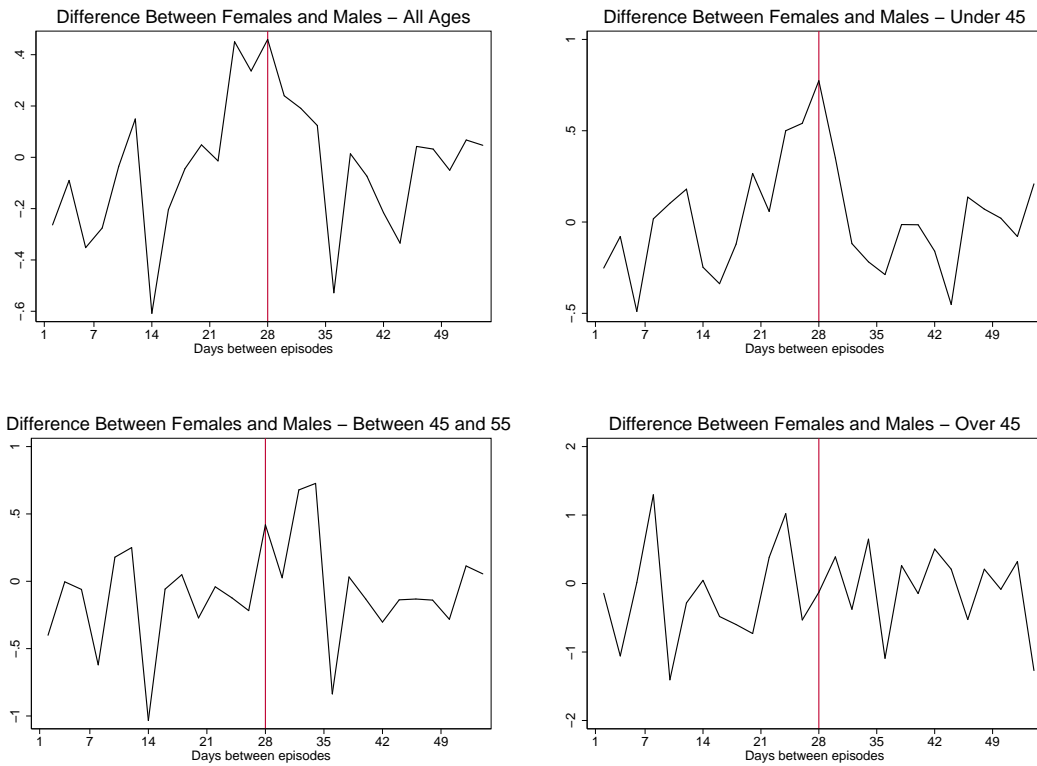
Notes: Standard errors in parenthesis. The estimated equation is equation 6. The excluded group is males with a small number of cyclical absences. All models control for the number of non-cyclical absences and for age. The observed gender gap is the coefficient on the female dummy in a regression that includes controls for the number of non-cyclical absences and for age (see column 2 in Table 6). The counterfactual gender gap is defined in equation 4. Sample includes workers 45 or younger.

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



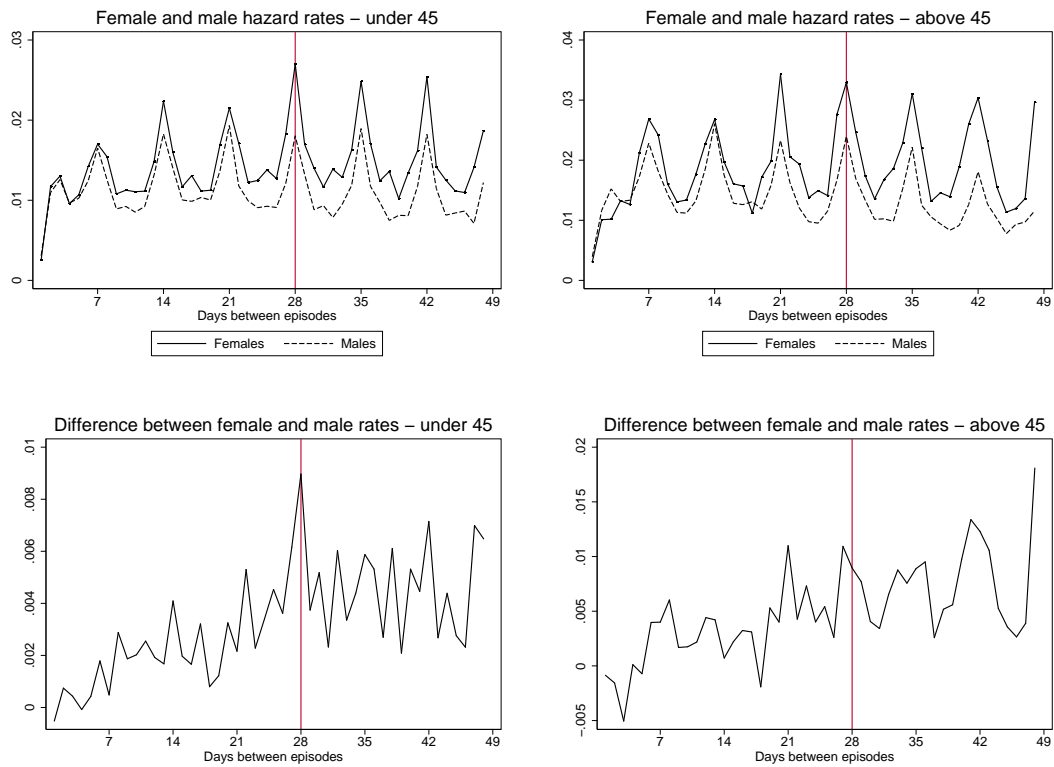
Note: The figure shows the female-male difference in the distribution of the number of days of absence between the beginning of two consecutive absence episodes.

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



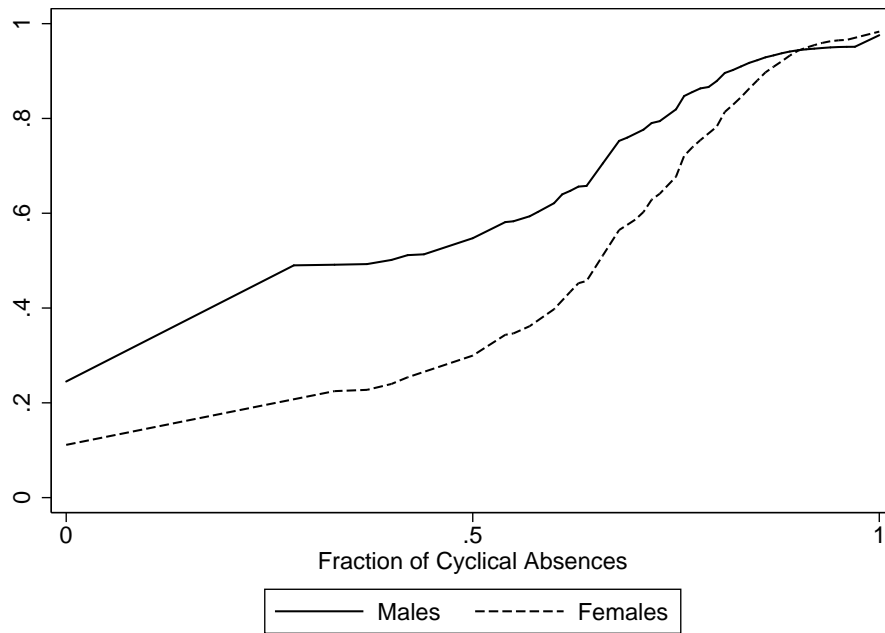
Note: The figure shows the female-male difference in the distribution of the number of days between the beginning of two absence episodes, calculated for all possible pairs of absences.

Figure 3: Hazard Rates, by Gender and Age



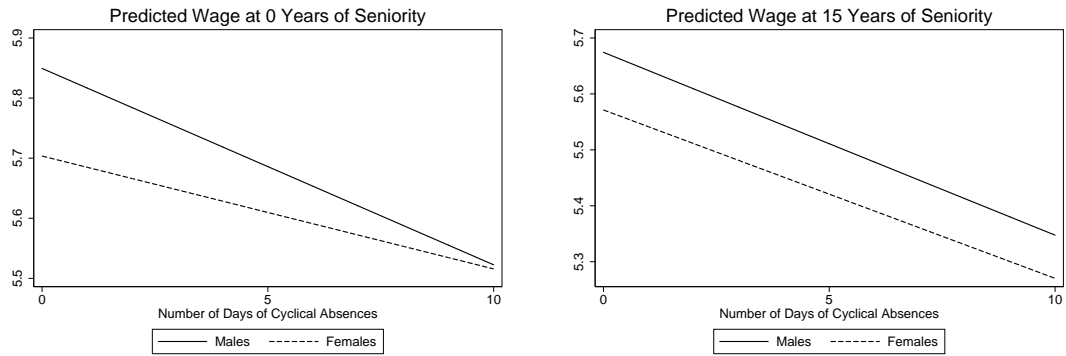
Note: The top panels show the Kaplan-Meier estimates of the hazard of an absence episode for male and females, with duration measured from the previous episode. The bottom panels show the female-male difference in the hazards.

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



Notes: The figure shows the cumulative distributions of the fraction of cyclical absences for males and females. The sample includes only workers 45 years old 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 7. 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: Standard errors in parenthesis. The sample includes full time workers continuously on 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.

Day at risk of absence	e^γ	Asymptotic t ratio	95% confidence interval		Log Likelihood
14	0.91	-1.34	0.80	1.04	-706736
21	0.82	-2.92	0.72	0.93	-706733
25	1.09	0.98	0.92	1.29	-706737
26	1.02	0.22	0.85	1.21	-706737
27	1.10	1.28	0.95	1.28	-706736
28	1.16	2.17	1.01	1.32	-706735
29	0.93	-0.86	0.80	1.09	-706737
30	1.16	1.66	0.97	1.38	-706736
31	0.91	-0.99	0.75	1.09	-706737
35	0.99	-0.11	0.86	1.14	-706737
42	1.06	0.81	0.91	1.22	-706737

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 days after a previous absence episode (see equation 1). The analysis is restricted to females younger than 45. The row for 28 corresponds to the second row and second column of Table 2.