TIME USE AND LABOR PRODUCTIVITY: THE RETURNS TO SLEEP

Matthew Gibson and Jeffrey Shrader*

Abstract—We investigate how the largest use of time—sleep—affects productivity. Time use data from the United States allow us to test a model in which sleep improves productivity. Consistent with theory, we find sleep is more complementary to home production than to leisure for nonemployed individuals. We then show that later sunset time reduces worker sleep and earnings. After ruling out alternative hypotheses, we implement an instrumental variables specification that provides causal estimates of the impact of sleep on earnings. A 1-hour increase in location-average weekly sleep increases earnings by 1.1% in the short run and 5% in the long run.

I. Introduction

QUESTIONS of labor productivity are fundamental to economics, important for both individual decisions and public policy. While there are traditions of research on the labor productivity effects of human capital (Becker, 1962, 1964) and health (Leibenstein, 1957; Mushkin, 1962), less attention has been paid to time use. Many types of time use, from reading to vacationing, plausibly have an impact on labor productivity in both market and nonmarket work. In this study, we examine the effects of the most common human time use: sleep.

Medical research indicates that sleep may play an important part in determining labor productivity. Tired doctors make more mistakes (Ulmer et al., 2009). Tired students perform worse on tests (Taras & Potts-Datema, 2005). Poor sleep impairs health (Cappuccio et al., 2010). Moreover for the average individual, sleep takes up more time than any other activity. Despite the manifest importance of sleep, economists have largely treated it as a biological phenomenon outside their purview. We investigate two important questions that have been overlooked almost entirely. First, how do individuals trade off sleep against other time uses? Second, how does sleep affect earnings and wages?

To examine the first question, we extend the time use model of Gronau (1977) to include productivity-enhancing sleep as a fourth time use in addition to labor, leisure, and home production. For workers, the model shows that if sleep improves wages more than it complements home production, then sleep and market work will be complements while sleep and home production will be substitutes. In the short run, there is evidence that this complementarity pattern holds, suggesting that sleep is relatively more beneficial for market work than home work among employed individuals. The model also allows us to test the productive sleep hypothesis on the sample of nonworkers. If an individual does not work, the model predicts that sleep and home production will be complements and sleep and leisure time will be substitutes. We find support for this hypothesis in both the short and long runs.

Our model implies that answering our second question is difficult because the relationship between sleep and wages is causal in both directions. Sleep increases labor productivity and higher wages raise the opportunity cost of sleep time. There is empirical evidence of the latter relationship—a pioneering study by Biddle and Hamermesh (1990) finds that higher wages are associated with less sleep—but the former relationship has not been studied. In addition, sleep may be correlated with unobservable worker characteristics that also influence wages.

Motivated by medical research on circadian rhythm, we resolve the endogeneity by using sunset time as a source of exogenous variation. In general the first-stage relationship is straightforward: earlier sunset causes workers to begin sleeping earlier, and because work and school start times do not respond as strongly to solar cues (Hamermesh, Frazis, & Stewart, 2008), this earlier bedtime translates into more sleep. In fact, sunset timing provides two types of variation: short run and long run. In the short run, within a location, earlier sunset in winter induces longer sleep duration. In the long run, comparing two locations in the same time zone, the location farther east will experience earlier average sunset than the location farther west. As a consequence, residents of the eastern location will sleep longer. These two types of sunset variation provide two instruments for sleep.

To implement our empirical strategy, we geocode observations from the American Time Use Survey (ATUS), which provides rich labor market information about individuals, a wealth of control variables, and detailed time use data from daily diaries. Using the diary date and location, we assign each observation a diary date sunset time and an annual average sunset time. We then use these variables to estimate the short- and long-run effects of sunset time on sleep and earnings, controlling—in the case of the short-run estimates—for fixed location characteristics, year effects, and individual characteristics, and—in the case of the long-run estimates—geographic characteristics (distance to the coast, latitude) and location-level demographic characteristics.

Consistent with our hypothesized first-stage relationship, we find that later sunset time significantly reduces sleep duration. The reduced-form effect of sunset time on earnings is also consistent with our sleep hypothesis. Intraannually, a 1-hour increase in sunset time decreases worker earnings by

Received for publication August 1, 2016. Revision accepted for publication November 27, 2017. Editor: Brigitte C. Madrian.

* Shrader: Columbia University and New York University; Gibson: Williams College.

We thank Prashant Bharadwaj, Julie Cullen, Gordon Dahl, Joshua Graff Zivin, Daniel Hamermesh, Erik Hurst, Mark Jacobsen, Craig McIntosh, Sendhil Mullainathan, Dale Squires, Michelle White, and seminar participants at UC San Diego for helpful comments. J.S. gratefully acknowledges support from the National Science Foundation Graduate Research Fellowship under grant DGE-1144086. The opinions, findings, and conclusions or recommendations expressed in this material are our own and do not necessarily reflect the views of the National Science Foundation.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00746.

The Review of Economics and Statistics, December 2018, 100(5): 783–798

doi:10.1162/rest_a_00746

0.4%, while a 1-hour difference in long-run average sunset time decreases worker earnings by 4.5%. We find these earnings changes stem almost entirely from wage changes rather than changes in hours worked. If labor markets are competitive and workers are paid their marginal revenue product, these wage changes derive from productivity changes. Using alternative econometric specifications, we rule out a number of other, nonsleep hypotheses and obtain similar results from multiple data sets in addition to ATUS. These results suggest that the exclusion restriction required for instrumental variables estimates—that the effect of sunset time on earnings operates only through sleep, conditional on our control variables—is reasonable. Under this assumption, we exploit sunset-induced sleep changes to identify both short- and long-run earnings effects.

Our results show that a short-run, 1-hour increase in location-average weekly sleep increases worker earnings by 1.1%. A permanent 1-hour increase in weekly average sleep increases location average earnings by 5%. These are, to our knowledge, the first causal estimates of how sleep affects earnings. Because our identification relies on location-level variation, these estimates should not be interpreted as individual effects. Both short- and long-run estimates potentially include productivity spillovers across workers. In addition, our long-run estimate may include general equilibrium effects induced by exogenously higher worker productivity. We find no evidence of nonlinearity in the sleep-earnings relationship.

Our study demonstrates that sleep is not just an economic curiosity; rather, it is a vital determinant of productivity and an important part of an individual's time allocation decision. A 1-hour increase in a location's weekly mean sleep raises earnings by roughly half as much as a one-year increase in education (Psacharopoulos & Patrinos, 2004).[1] These results point to the large impact that non-labor-market activities can have on labor market performance. By examining the largest use of human time, our study contributes to the time use literature following Becker (1965). It complements the important work on the evolution of leisure time by Aguiar and Hurst (2007). Our study also contributes to the growing literature on how environmental forces influence worker productivity (Graff Zivin & Neidell, 2012) and to the broader productivity literature on factors like information technology (Bloom, Sadum, & Van Reenen, 2012) and workplace practices (Black & Lynch, 2001).

The rest of the paper proceeds as follows. Section II presents a time use model with sleep as a choice variable, illustrating identification challenges, and discusses related literature. Section III presents the estimating equations and discusses our identification strategy. Section IV describes the data used in the study. Section V reports the main results. First, we test model predictions about trade-offs across sleep and other time uses. We then report the reduced-form and first-stage effects of sunset time on earnings and sleep, followed by an extensive set of robustness checks aimed at instrument validity. Finally we present instrumental variables estimates of the effect of sleep on earnings. Section VI concludes.

[1] While the effect of increasing earnings for all workers in a location might differ from the partial equilibrium estimate of Psacharopoulos and Patrinos (2004), the latter nonetheless provides an instructive benchmark.

II. Identifying the Effect of Sleep on Productivity, Wages, and Earnings

A. *Previous Research*

Existing studies of the relationship between sleep and wages in economics are few and largely concerned with addressing the question of whether sleep should be treated as a choice variable rather than simply a biological necessity. Biddle and Hamermesh (1990) were the first to provide empirical evidence on this issue, and their paper remains one of the only empirical investigations of labor market impacts of sleep. They lay out a model with agents optimizing over sleep, work, and leisure time in an otherwise standard setting. While their theoretical model allows sleep to affect productivity, Biddle and Hamermesh do not focus on this relationship in their empirical work. Instead they emphasize the causal mechanism operating in the opposite direction, modeling sleep as a function of instrumented wage (see, e.g., Biddle & Hamermesh, 1990, table 6). Brochu, Armstrong, and Morin (2012) and Szalontai (2006) also estimate the impact of wage changes on sleep using more recent data from Canada and South Africa. Finally, Bonke (2012) has examined the impact of two chronotypes—whether the individual is a "morning" or "evening" person—on income, finding that morning types earn more.

Daylight savings time (DST) has been used in a variety of settings in economics as a proxy for sleep changes. For example, Smith (2016) finds that the spring DST transition results in more automobile accidents and attributes the change to sleepiness behind the wheel. However, the short-term nature of any sleep change induced by DST limits its use in studying slower-moving outcomes like wages. Moreover, while the spring transition into DST reduces sleep by 40 minutes on the day of the change, the transition out of DST does not induce a change in sleep (Barnes & Wagner, 2009).

Among medical studies, Van Dongen et al. (2003) conducted a two-week, laboratory-controlled study on the relationship between sleep levels and cognitive performance. The researchers placed subjects into groups receiving 4, 6, and 8 hours of sleep. The subjects were given daily tests of attention, memory, and cognition. The research found that relative to the 8-hour group, the groups subjected to 4 and 6 hours of sleep performed progressively worse on all three tests. Moreover, the performance decline was linear across the treatments. Intriguingly, the subjects' subjective assessments followed a different pattern, declining for a few days and then leveling off. Observed cumulative effects quickly achieved large magnitudes. After one week, subjects in the 6-hour group performed as badly as subjects who were deprived of sleep entirely for one night. This indicates that sleep reductions well within the observed range, continued over long periods of time, can have very large effects.

B. A Productive Sleep Model

A formal treatment of time use in the presence of productivity-enhancing sleep clarifies the identification challenge stemming from reverse causality and generates predictions about time use trade-offs for workers—a group for whom we can observe wages or earnings as a proxy for productivity—and nonworking individuals.[2] These predictions, along with similar predictions for workers, can be used to test whether sleep is productive. The model shows that productive sleep implies that sleep and work (either in the market or at home) can be complements. In contrast, if sleep were not productive, it would be a substitute with the other time uses. We empirically investigate trade-offs between time uses in section VA and find that sleep is indeed complementary to labor and home production.

To model productive sleep, we present an extension of the time use model of Gronau (1977). The individual's problem is to maximize a utility function $u(Z(x, T_L))$ where Z is a home production function in the style of Becker (1965) that takes goods, x, and leisure time, T_L, to produce consumables. Assuming as usual that u is increasing and quasi-concave, we can ignore the utility function and consider the individual's problem to be one of maximizing Z subject to the constraints below.

Goods can either be purchased in the market, at a price normalized to 1, or produced at home. Denote market goods as x_M and home-produced goods as x_H so that total goods are given by $x = x_M + x_H$. Work time is denoted T_N, and the individual can gain market goods by working at wage $W(T_S) = \alpha + w(T_S(\theta))$, which is a function of an exogenous parameter, α, and time spent sleeping, $T_S(\theta)$. Sleep time is chosen by the individual, but it is influenced by an exogenous price, denoted θ, which one can think of as sunset time. We assume that a higher price (or later sunset time) reduces sleep time, so $dT_S(\theta)/d\theta < 0$. Since we are modeling sleep as productivity enhancing, we also assume that more sleep will, ceteris paribus, increase wage, so a change in wage due to an exogenous change in θ, $\frac{dW}{d\theta} = \frac{\partial W(T_S)}{\partial T_S}\frac{dT_S(\theta)}{d\theta} := w'(T_S)T_S'(\theta)$, will be negative. For convenience, T_S will be written with θ suppressed unless we explicitly want to model a change in sleep time due to an exogenous price change. The individual has nonlabor income V so that $x_M = W(T_S)T_N + V$.

Home goods, x_H, are produced using the production function $x_H = f(T_H, T_S)$, where T_H is home production time. Given the assumption that sleep is work productive, we naturally also assume that the change in home production with

[2] In general, earnings reflect both work hours and productivity. In what follows, we provide evidence that the earnings changes we estimate stem almost entirely from wage changes rather than hours changes. This need not be true in other settings.

FIGURE 1.—PRODUCTIVE SLEEP

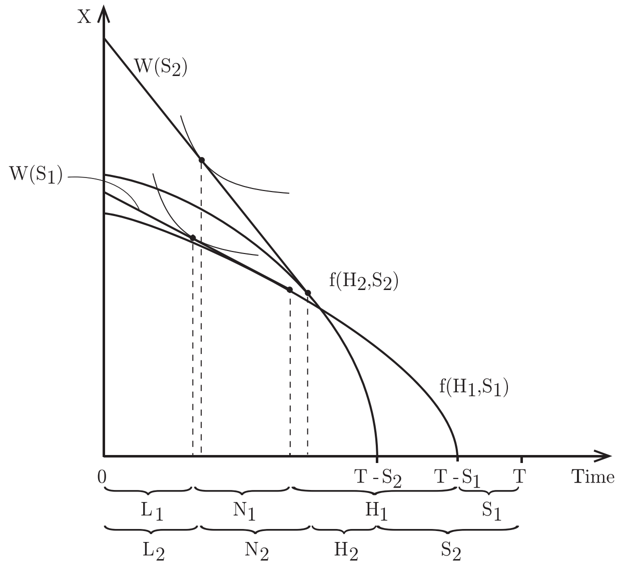

A graphical depiction of the productive sleep model in the style of Gronau (1977, figure 1). The figure depicts two possible choices for levels of sleep, showing that increasing sleep raises both work and home productivity, at the cost of taking time away from the total time budget. Time uses are denoted by their subscripts in the model for legibility.

respect to sleep, f_2', is positive. Assume $f(0, .) = 0$ so that $T_H = 0$ is equivalent to $x_H = 0$.

Putting all time uses together, the total time constraint is $T = T_L + T_H + T_N + T_S$. Finally, assume that sleep cannot be used as leisure to produce goods in Z.[3] Substituting the time budget into the goods budget, the combined budget constraint is

$$x_M + W(T_S)(T_H + T_L + T_S) = W(T_S)T + V \tag{1}$$

and the optimization problem is

$$\max_{T_L, T_H, T_S, x_M} \; Z\left(x_M + f\left(T_H, T_S\right), T_L\right) \tag{2}$$

$$+ \lambda_1\big(W(T_S)T + V - x_M$$

$$- W(T_S)\left(T_H + T_L + T_S\right)\big) \tag{3}$$

$$+ \lambda_2 x_M + \lambda_3 T_H + \lambda_4 T_S. \tag{4}$$

This is a concave problem since $W(T_S)T > W(T_S)(T_H + T_L + T_S)$.

The main predictions of the model relative to the case where sleep is not productive are, first, that sleep can be complementary to either home production or labor for employed individuals and, second, that sleep can be complementary to home production for individuals who do not work.

We can get intuition for these results by utilizing Gronau's graphical device of representing all activities in just two dimensions: goods and net-of-sleep time. In figure 1, the home production function, $f(T_H, T_S)$, should be viewed as a projection onto the $(x_H, -T_H)$ plane for a given value of T_S. The first-order conditions show that the wage function will

[3] We discuss relaxation of this assumption at the end of this section.

be tangent to the home production function. This point of tangency is where a worker will choose to divide time between leisure and work on the left and sleep and home production on the right. The tangency between utility and wages will determine the division between leisure and market work.

Changes in sleep influence the slope of the wage line, the steepness of the home production function, and the amount of time available for other time uses. The figure depicts two possible values for sleep: $T_{S,1}$, a low value, and $T_{S,2}$, a high value. By assumption, an increase in sleep raises the slope of the wage line. If this were the only effect, the wage, $W = \alpha + w(T_{S,2})$, would be tangent to the home production function $f(T_{H,1}, T_{S,1})$ at a point farther to the right, causing the worker to expand total time devoted to work and leisure. Sleep also raises the productivity of home production, however, causing a reduction in time devoted to leisure plus work. For the changes shown in the figure, the wage increase from sleep is large enough to lead to an overall increase in the amount of market work and a decrease in the amount of home production. This provides intuition that in cases where sleep is especially beneficial for productivity in market work, sleep and labor will be complements. A similar condition holds for home production, as the following results make clear.

To formalize the time use predictions, consider an individual who works in the market and at home. This case corresponds to ($x_M > 0, T_H > 0$), so both λ_2 and λ_3 are 0. Assuming that sleep time is positive, $f_1' > 0, f_2' > 0, f_{11}'' < 0$, and $f_{22}'' < 0$, the first-order conditions can be written as

$$\frac{Z_2'(x, T_L)}{Z_1'(x, T_L)} = W(T_S), \tag{5}$$

$$f_1'(T_H, T_S) = W(T_S), \tag{6}$$

$$f_2'(T_H, T_S) + T_N w'(T_S) = W(T_S), \tag{7}$$

plus budget constraints. These equations implicitly define optimal time use quantities in terms of parameters, including sunset time, θ.

Taking a total derivative of equation (6) with respect to θ and rearranging yields

$$\frac{dT_H}{d\theta} = \left(\frac{f_{12}''(T_H, T_S) - w'(T_S)}{f_{11}''(T_H, T_S)}\right)\frac{dT_S}{d\theta}. \tag{8}$$

By assumption, $\frac{dT_S}{d\theta} < 0$, so the right-hand side will be positive if $w' > f_{12}''$. In words, if the marginal increase in wages with respect to sleep is larger than the complementarity between sleep and home production, then home production will rise in response to a later sunset time. If the opposite condition holds, then later sunset time will decrease both sleep and home production time.

To see that sleep and work can also covary in the model, consider equation (7). Again take the total derivative with respect to θ to find

$$\frac{dT_N}{d\theta} = \left(w'(T_S)\right)^{-1} \times \tag{9}$$

$$\left(\left(w'(T_S) - f_{22}''(T_H, T_S) - T_N w''(T_S)\right) \times \frac{dT_S}{d\theta} - f_{12}''(T_H, T_S)\frac{dT_H}{d\theta}\right). \tag{10}$$

Assuming that $W_{11} \leq 0$, if $\frac{dT_H}{d\theta} > 0$ and $f_{12}'' > 0$, then the whole term will be negative.[4] Therefore, a later sunset time would lead to less sleep and less work. So for a worker, depending on f_{12}'' and w', productive sleep can be complementary to either home production or work time.

For an individual who does not work outside the home, the predictions are starker than for a worker: as long as the labor productivity gain from more sleep is not so large that it induces the person to work, then home production should increase with an increase in sleep and leisure time should fall. The formal statements for these relationships can be derived from the first-order conditions, taking $\lambda_3 = 0$ and $\lambda_2 > 0$. In particular, in this model, an individual chooses not to work if the derivative of the home production function with respect to home production time is steeper than the market wage rate. At the margin, a change in sleep will not change this condition, so it must be the case that sleep and home production will be complements by the same logic used to sign equation (8) above.

Contrast these predictions with a model of sleep that is not productivity enhancing. Suppose that sleep is just a biological necessity: it grants neither utility nor productivity. In such a case, sleep only removes time that could have been devoted to other tasks. A later sunset time will decrease sleep, causing a pure income effect for the individual. Leisure, work, and home production time will rise, with the exact split between these time uses dependent on the production function, Z. Therefore, sleep will be a substitute for all other time uses.

If sleep time can be used for consuming goods but does not improve productivity, it is indistinguishable from leisure. Therefore, sleep and leisure will be substitutes, and from figure 1, one can see that inducing a worker to consume more leisure will decrease work time but leave home production time unchanged. Therefore, such a model would predict that sleep and work would be substitutes while sleep and home production would be unrelated (at the margin). Neither of these models generates the prediction that sleep will be complementary to either work or home production.

Finally, we can use the model to motivate our empirical method by highlighting the potential reverse causality stemming from an increase in wages. Applying the implicit

[4] Medical studies of the effect of sleep on health often find a nonlinear relationship (e.g., Cappuccio et al., 2010) that suggests $W_{11} < 0$). Assuming $W_{11} = 0$ might be reasonable, at least for some range of sleep hours. Van Dongen et al. (2003), for instance, finds that performance on attention tasks declines linearly with sleep deprivation of up to 4 hours per night. To the limited extent that we can test for it in our empirical setting, we do not find evidence for nonlinearity in the sleep-wage relationship.

function theorem to the first-order conditions, one can formally find the effect of an exogenous change in α on sleep. The expression for this change can be found in the online appendix. Intuitively, if sleep is relatively more important for work than for home production, then a dominant income effect coming from increased wages can cause the agent to reduce sleep in exchange for more home production. A naive regression of wages on sleep would yield a negative coefficient in this case even if sleep were productive.

III. Empirical Strategy

A. Estimating Equations

We would like to recover the relationship between sleep and wages, where $\partial W/\partial T_S > 0$ would provide evidence for productivity-enhancing sleep. Given the reverse causality between wages and sleep, however, we might erroneously find $\partial W/\partial T_S < 0$. To avoid this problem and account for the wide variety of other omitted variables that might covary with sleep and wages, we predict sleep using two instruments based on local sunset time and then use the instrumented values of sleep to estimate earnings and wage impacts.

The first instrument uses daily variation in sunset within a given location. Because this instrument varies on a daily basis, it will identify short-run variation in sleep (Frazis & Stewart, 2012). Thus, we can use it to estimate a short-run first stage,

$$T_{S,ijt} = \alpha_1 \text{sunset}_{jt} + \gamma_{1,j} + \mathbf{x}_{it}'\delta_1 + \eta_{1,ijt}, \tag{11}$$

and reduced form,

$$\ln(W_{ijt}) = \alpha_2 \text{sunset}_{jt} + \gamma_{2,j} + \mathbf{x}_{it}'\delta_2 + \eta_{2,ijt}, \tag{12}$$

where $T_{S,ijt}$ is nighttime sleep for individual i in location j on date t, sunset_{jt} is the sunset time on that date in that location, γ_j is a location fixed effect, $\mathbf{x}_{it}$ is a vector of individual level controls, W_{ijt} is a measure of wages or earnings observed at time t, $\eta_{1,ijt}$ is the error term for the first stage, and $\eta_{2,ijt}$ is the error term for the reduced form. Controls are four race indicators; age and age squared; a gender indicator; indicators for holidays, day of week, and year; and detailed occupation code indicators. More details on these control variables can be found in the data discussion in section IV. Following the suggestions of Winship and Radbill (1994) and Solon, Haider, and Wooldridge (2015), we do not weight observations, but we do control for weekends since they are oversampled in our data set.

If seasonal sunset time is a valid instrument for sleep, then this first stage and reduced form can be used to construct causal estimates of the effect of sleep on earnings or wages by taking the ratio of α_2 to α_1. In practice, we will calculate the instrumental variables estimator using two-stage least squares. We denote this estimate β_{SR}, or the effect of sleep in the short run.

The second instrument is annual average sunset. This instrument exploits spatial differences in sunset time within and across time zones. Because this is a fixed feature of a location, it will identify long-run differences in sleep (Frazis & Stewart, 2012). For estimation, we collapse the ATUS data to the location level. This serves to emphasize that variation in long-run sunset time is permanent and common to all workers in a location. We then estimate the following first stage,

$$T_{S,j} = \varphi_1 \text{sunset}_j + \mathbf{x}_j'\zeta_1 + \varepsilon_{1,j}, \tag{13}$$

and reduced form,

$$\ln(W_j) = \varphi_2 \text{sunset}_j + \mathbf{x}_j'\zeta_2 + \varepsilon_{2,j}, \tag{14}$$

where $T_{S,j}$ is average nighttime sleep in location j, sunset_j is the average sunset time in that location, $\mathbf{x}_j$ is a vector of controls, W_j is average earnings or wage in that location, and $\varepsilon_{k,j}$ is an error term for $k \in \{1, 2\}$. We control for both geographic characteristics (coastal distance and latitude) and demographics (gender, age, race and occupation shares, plus population density).

Following the recommendation in Solon et al. (2015), we weight location-level observations using counts of the underlying individual ATUS observations to correct for heteroskedasticity. (The online appendix provides evidence of heteroskedasticity from a modified Breusch-Pagan test, presents unweighted results, and explores differences between weighted and unweighted results.)

Again, if average sunset time is a valid instrument, the causal effect of long-run changes in sleep can be found by taking the ratio of φ_2 to φ_1. We denote this coefficient β_{LR} and estimate it by two-stage least squares. Although the control variables primarily serve to reduce residual variance, as discussed in section V, we do find evidence that average sunset is not unconditionally exogenous with respect to coastal distance and population density (see the appendix). Therefore the identifying assumption underlying our long-run IV estimates is one of conditional exogeneity, as discussed in sections IIIB and VC.

B. Local Sunset Time Instruments

Relevance of Sunset to Sleep. The relevance of sunset time as an instrument for sleep stems from the biological relationship between sleep patterns and daylight. Human circadian rhythm is synchronized with the rising and setting of the sun through a process known as entrainment. This force is powerful, with Roenneberg, Kumar, and Merrow (2007) showing that "the human circadian clock is predominantly entrained by sun time rather than by social time." Using data from Germany, the authors demonstrate that living in a location with a later sunset induces individuals to begin sleep later. The detailed ATUS files enable us to reproduce this result: workers experiencing a later sunset go to bed later, and this causal connection between sunset and bedtime persists even if the worker goes to bed well after dark. Intra-annual changes in sunlight also influence human sleep patterns through a similar process of entrainment (Hubert, Dumont, & Paquet, 1998). In a vacuum, a later sunset time might cause workers to go to bed later and also rise later, leaving sleep duration unchanged. But workers face morning coordination constraints due to work and school scheduling (Hamermesh, Myers, & Pocock, 2008), so later sunset and later bedtime decrease sleep duration, both intra-annually and in the long run. Our estimates in section VB are consistent with these hypotheses. The evidence bearing on instrument validity is different for the short and long run. We discuss the two cases separately.

Short-run validity. For our short-run estimates of the effect of sleep on wages and earnings, validity requires that other wage determinants not covary with daily sunset time within a location. The primary threat to this assumption is seasonally varying wage determinants, since sunset time exhibits a regular seasonal pattern, as shown in figure 2a for all ATUS observations. One can see that sunset time is generally described by a cosine wave with a period of one year. This wave is phase-shifted by roughly ten days relative to the calendar year. The amplitude of the wave is determined by the latitude of the location, and vertical translations are due to within-time zone variation, which we use for our long-run estimate. The final important features of sunset time are the prominent jumps in the spring and fall caused by daylight savings time. A 2005 change in the timing of daylight savings observance makes these jumps nonsharp.[5]

Knowing the exact pattern followed by sunset time allows us to characterize the degree of potential confounding from other seasonal variables. These calculations are given in section VC. Because of daylight savings time and the phase shift relative to the calendar, our estimates are robust to a wide range of seasonal confounders. Moreover, these features allow us to clearly disentangle short-run variation in sunset from calendar features like the December shopping season.

Our productivity measures are usual weekly earnings and usual hourly wage rather than wages on the day of the interview. This raises an additional identification issue inherent in studying earnings and wages rather than productivity: timing mismatch between observations of sleep and wages combined with a potentially low-frequency relationship between wages and productivity. These issues mean that our short-run estimates of both the reduced-form equation (12) and the resulting two-stage least squares estimator will be attenuated. In the appendix, we derive an expression for asymptotic bias that depends only on the frequency of wage or earnings changes and a trigonometric function derived from solar mechanics. Using that expression we bound the degree of attenuation between 0.25 and 1, so we expect a priori that

[5] For more discussion of solar mechanics, see appendix section 2.

FIGURE 2.— SUNSET VARIATION IN THE SHORT AND LONG RUNS

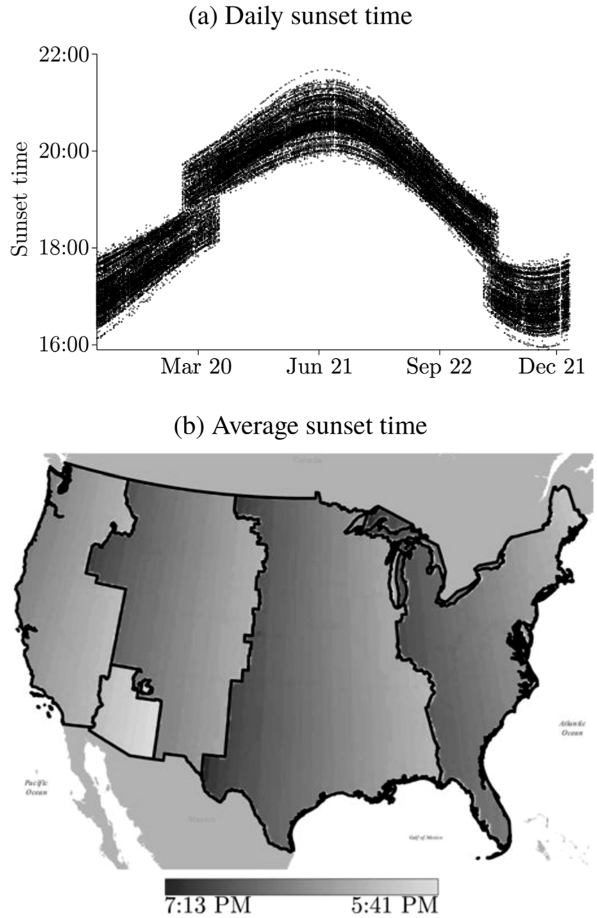

(a) The vertical axis is measured in 24 hour time (e.g., 19:00 is equivalent to 7:00 p.m.). Each point shows the sunset time for a location sampled by ATUS within the continental United States. (b) The map shows sunset time at the vernal equinox for the continental United States in 2012, which is a close approximation to average sunset time. Darker color indicates later sunset, and lighter color indicates earlier. The time zone boundaries are given by bold black lines.

our short-run estimates should be between 0 and one-quarter of the true parameter value.

There is one short-run identification issue we cannot address: seasonal variation in sunset time is almost perfectly correlated with seasonal variation in sunrise time and daylight duration.[6] Therefore, in purely statistical terms, all short-run results could be recast in terms of either of these other variables. Our interpretation of the short-run results could be incorrect if daylight affects both earnings and sleep through mood, the hedonic value of leisure, or some other channel.[7] Medical studies find positive effects of daylight duration on mood (Murase et al., 1995; Lambert et al., 2002), so if this is the dominant confounder, our earnings and wage estimates will be biased downward in magnitude. We focus on sunset time rather than sunrise time because it is emphasized by existing medical literature and because it is the

[6] Only changes in daylight savings time break this linkage.

[7] Daylight duration does not create an analogous problem for the long-run estimates, as all locations in the continental United States experience nearly the same average amount of daylight.

driver of long-run differences in sleep, as discussed in the next section.

Long-run validity. Figure 2b illustrates long-run variation in sunset time across locations. As the sun sets, eastern locations grow dark earlier than western locations, leading residents in more easterly locations to go to bed earlier and sleep longer. By design, the maximum difference in sunset time within a U.S. time zone is approximately 1 hour.

The difference in average sunset time between two locations over the year is plausibly unrelated to other factors influencing the labor market. In particular, time zone boundaries break the link between average sunset time and longitude. Average sunset time is also, by construction, orthogonal to latitude. All locations in the continental U.S. experience approximately the same average daylight duration over the year, so this is not an omitted variable in our long-run analysis, which is restricted to this geographic area. If average sunset time shifted the timing of work within the day and if workers were more productive at particular times of day, that could violate the exclusion restriction. Hamermesh et al. (2008), however, find that the response of work schedules to sunrise and sunset times on the day of an ATUS diary is extremely small and not statistically significant. This suggests work schedules do not adjust to sunset time in either the short or long run.

Railroads implemented the first U.S. time zones, called Standard Railroad Time (SRT), on November 18, 1883. They replaced a patchwork of railroad time standards and were quickly adopted by the U.S. government and Western Union (Allen, 1883; "The New Standard Time," 1883). While railroads were the first adopters, the primary impetus for standard time and the zone plan itself came from scientists concerned with problems like simultaneous observation of the aurora borealis at different points across the United States (Bartky, 1989). The width of a zone, 15 degrees of longitude, was chosen to correspond with a 1-hour difference in solar time (Library of Congress, 2010).

Endogenous modifications to time zone borders could have undermined this initial randomization. State and local governments may petition the Department of Transportation (DOT) to switch time zones. The DOT criterion for evaluating proposed time zone changes is "the convenience of commerce."[8] This process means that the precise location of the boundary is endogenous, and a regression discontinuity design comparing nearby communities on opposite sides of the boundary could be biased. Moreover, the integration of labor markets across time zone borders in some counties could pose identification problems. (In table 6 we show our results are robust to the exclusion of all counties on time zone borders.) To avoid potential endogeneity, we drop locations that do not observe daylight saving time. Finally, while time zone borders often coincide with state borders, they frequently do not, and 12 of the lower 48 U.S. states span multiple time zones (Hamermesh et al., 2008).

Current or past worker sorting on sunset time could threaten instrument validity if such sorting induced correlation between average sunset time and characteristics that influence worker productivity. We investigate this possibility in the appendix and find no evidence for it. Potentially more problematic in our finite sample is a spurious correlation between the instrument and location-level determinants of productivity. We do find a statistically significant relationship between average sunset time and population density, which motivates our use of a flexible control for this variable. Within a time zone, long-run average sunset time is a linear function of longitude and thus also correlated with coastal distance. Table 6 suggests coastal distance is indeed an important potential confounder, and we control flexibly for it in all our specifications. Conditional on our flexible controls, we find no evidence that early- and late-sunset workers differ on observables. The identifying assumption in our long-run models is exogeneity of average sunset time conditional on these controls. Having found two important variables correlated to our instrument, it is natural to ask whether there might be others. We cannot exclude this possibility. In section VC, we investigate a wide range of possible confounders and find our results are robust.

IV. Data

The largest data set containing both sleep time and wage information is the American Time Use Survey (ATUS), administered by the U.S. Bureau of Labor Statistics (BLS) since 2003. ATUS includes a random sample drawn from households that have recently completed participation in the Current Population Survey (CPS). For example, if a household participated in the CPS from January through April 2013 and January through April 2014, it would be eligible for ATUS sampling in June, July, or August 2014.

The BLS collects the data using computer-assisted telephone interviews, which cover the period from 4:00 a.m. on the previous day to 4:00 a.m. on the interview day. By employing a very short recall period and forcing all time uses to sum to 24 hours, this method minimizes the possible biases associated with time diaries (Hamermesh et al., 2005). Interviewers use conversational techniques that have been shown to reduce recall bias in laboratory settings (Schober & Conrad, 1997). For each time use, the interviewer records either the duration or the end time. In most cases, the interviewer records the respondent's description verbatim, and it is coded later. Codes are very detailed, distinguishing, for example, between watching live games of hockey and basketball. Because the sample is drawn from CPS respondents, CPS demographic and labor variables are available for almost all ATUS respondents. Each respondent participates in ATUS only once, however, so it is not possible to construct an individual time use panel. (For a more detailed description of ATUS, see Hamermesh et al., 2005, and Bureau of

[8] For details on the DOT process, see Valpando (2013) and the appendix.

Labor Statistics, 2015.) For the analysis of employed individuals, we use the sample of prime-age individuals who report receiving positive weekly earnings from a primary or secondary job and who work full time.[9] (Summary statistics for the employed sample, broken down by sunset time, are given in table 4.)

To assign locations to individuals in ATUS, we first merge ATUS data with the corresponding CPS data. For a given individual, the CPS data often contain location at the county level. This variable is censored for individuals living in counties with fewer than 100,000 residents. When county is available, we assign the county centroid as an individual's location. We have county location for approximately 44% of ATUS observations. For an additional 28% of observations, we observe location at the level of Census CBSA, a small group of counties in the same metropolitan area. In total, we are able to geocode 72% of observations at the substate level.[10] For the remaining 28% of observations, ATUS contains location at the state level. We assign the 2010 population-weighted state centroid (computed by the Census) as the location for these individuals. Where we refer to Federal Information Processing Standards (FIPS) codes, we are referring to either the county (FIPS 6-4) or CBSA-level code, if available, or the state-level code (FIPS 5-2) where more detailed location is unavailable. Using the interview date and respondent location, we are able to determine sunset time for each individual in the data set using solar mechanics algorithms from Meeus (1991). We compute annual average sunset time by computing sunset for each day in an individual's location, then calculating the mean over days of the year.

Our sleep variable is nighttime duration from the ATUS diary, multiplied by seven to obtain a weekly measure. We employ a weekly rather than a daily sleep variable to match the frequency of our earnings variable. We remove any sleep that starts and ends during daylight hours on the date of diary entry. This will exclude naps, which might be an adaptation strategy for some short sleepers; however, it also removes night-shift workers, for whom the sunset instrument should not be relevant. Empirically, our point estimates are practically unchanged by the exclusion of daytime sleep, but precision of the first stage increases substantially.[11]

Our primary labor productivity measure is "usual weekly earnings" as reported in ATUS. While this measure could capture changes in both wages and hours worked, our earnings results largely reflect wage changes, as discussed in sections VA and VC. Work time responds less than sleep time to sunset, and explicitly holding work hours fixed does not substantially change our estimates. The usual weekly earnings variable is defined for all respondents who have positive labor income and are not self-employed. It is top-coded above $2,884.61. We also estimate a version of our model including only workers who receive an hourly wage—"hourly earnings at main job"—as reported in ATUS. This variable is also top-coded at the level such that hourly earnings multiplied by usual weekly hours equal $2,884.61. Among employed individuals in the primary estimation sample, 52% of individuals report an hourly wage. Some control variables (e.g., occupation codes) appear in both ATUS and CPS files. Where possible, we use ATUS variables, which are more recent. Our preferred regression specifications include a set of 22 occupation dummies or shares based on the ATUS trdtocc1 variable, which categorizes the respondent's main job. Examples include "education, training, and library occupations" and "food preparation and serving related occupations." As discussed in section IIIA, in our long-run analysis, we collapse ATUS to a cross section in locations.

To investigate the robustness of our long-run results, we also use the Quarterly Census of Employment and Wages (QCEW), which covers all U.S. counties. The QCEW is collected by the BLS as part of the Covered Employment and Wages Program (ES-202), under which state employment security agencies report data to BLS. Covering all workers eligible for unemployment insurance, the data provide "a virtual census" of nonagricultural employees and cover 47% of agricultural employees (Bureau of Labor Statistics, 1997). BLS estimates that the QCEW covers 96% of civilian jobs, accounting for 92.5% of the wage and salary component of national income. Self-employed workers and members of the U.S. military are excluded. Earnings measures include bonuses, stock options, and gratuities but exclude employer contributions to unemployment insurance, worker's compensation, and pensions (Bureau of Labor Statistics, 1997). We collapse a county-level panel (1990–2013) to a cross section in order to investigate the reduced-form effects of our long-run instrument. The appendix presents summary statistics.

V. Empirical Results

A. *Time Use Complementarity*

We begin by testing theoretical predications about time use trade-offs. Our model predicts that if sleep is relatively more work-productivity-increasing than home-productivity-increasing, an increase in sleep should increase work time and decrease home production time. Moreover, leisure should respond negatively to an increase in sleep. For nonworkers, sleep should be complementary to home production.

[9] We exclude individuals in locations that do not observe daylight savings time.

[10] Even for county-geocoded observations, there will be measurement error in location and therefore sunset time. Such error arises because the population-weighted county centroid in the ATUS sample may differ from the spatial centroid. Under plausible classical assumptions, this will attenuate both first-stage and reduced-form estimates but will not attenuate IV estimates. County size is roughly constant within a time zone, but greater in the Pacific and Mountain zones than in the Central and Eastern zones.

[11] ATUS gathers data on all sleep during the course of a single 24-hour period for each individual, so there are potentially other ways to calculate naps, and our results are robust to alternative definitions.

TABLE 1.—TIME USE RESPONSES TO SUNSET: SHORT RUN

	Sleep	Work	Leisure	Home Production
Full-time workers				
Daily sunset time	−0.38***	−0.11*	0.21***	0.085
	(0.041)	(0.065)	(0.060)	(0.055)
Mean of dependent variable	57.9	31.2	49.5	26.6
Observations	61,161	61,161	61,161	61,161
Nonworkers				
Daily sunset time	−0.38***		0.26***	−0.19*
	(0.065)		(0.097)	(0.10)
Mean of dependent variable	62.5		61.5	38.4
Observations	28,126		28,126	28,126

The table shows results from seven separate regressions estimating versions of the first-stage equation (11) on ATUS data, where the dependent variable is time use in one of four categories, indicated at the top of each column. Controls, number of observations, and standard error clustering are the same as in table 3. The sample of full-time workers is the same as that used throughout the paper. The nonworkers are prime-age individuals who do not report having a full- or part-time job. Significance indicated by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.1$.

TABLE 2.—TIME USE RESPONSES TO SUNSET: LONG RUN

	Sleep	Work	Leisure	Home Production
Full-time workers				
Average sunset time	−0.93***	0.66	0.73	−0.28
	(0.28)	(0.62)	(0.57)	(0.47)
Mean of dependent variable	57.9	31.4	49.4	26.5
Observations	529	529	529	529
Nonworkers				
Average sunset time	−0.11		1.49**	−1.30*
	(0.48)		(0.71)	(0.79)
Mean of dependent variable	62.5		61.5	38.1
Observations	527		527	527

The table shows results from seven separate regressions estimating versions of the first-stage equation (11) on ATUS data where the dependent variable is time use in one of four categories, indicated at the top of each column. Controls, number of observations, and standard error clustering are the same as in table 4. The sample of full-time workers is the same as that used throughout the paper. The nonworkers are prime-age individuals who do not report having a full- or part-time job. The weights in the long-run sample are based on the population of all workers in each location. Significance indicated by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.1$.

Tables 1 and 2 show estimates of these time complementarities for full-time workers and nonworkers in both the short-run and long-run ATUS samples. To estimate complementarity, we use a new method by Allen and Rehbeck (2016). With a simplex budget, as in the case of time use, complementarity can be assessed using a ratio estimator similar in spirit to a Wald estimator. In our setting, this means that we can determine whether two time uses are complements or substitutes by regressing each on sunset time, then taking the ratio of the two resulting regression coefficients. A positive ratio indicates that the time uses are complements, while a negative ratio indicates they are substitutes. This measure is theoretically motivated but also intuitively appealing: because the time budget must bind, any two large uses of time are likely to be negatively correlated. Therefore, complementarity must be assessed by comparing changes in the two time uses with respect to a time shifter.

To bring this method to the data, we categorize time uses from the ATUS diaries as either sleep, work, leisure, or home production.[12] Estimating the effect of sunset time on these four time use categories, one can see that they provide evidence in favor of the productive sleep model. In the short-run sample, complementarity between work and sleep for workers and complementarity between home production and sleep for nonworkers are consistent with productive sleep and are inconsistent with a model where sleep is not productive. Similar evidence is provided by the long-run relationship between sleep and home production in both the worker and nonworker sample, although the evidence is less precise for workers.

In the context of the model, the short-run worker results show that sleep is relatively more work-productivity-enhancing than home-productivity-enhancing because instrumented sleep and work positively covary while instrumented home production and sleep negatively covary. For the long-run, full-time worker sample, there is not a precise relationship between instrumented sleep and any of the other three time categories. The point estimates indicate that sleep is relatively more home-productivity-enhancing for both the worker and nonworker samples.[13]

The other results do not help distinguish between productive and nonproductive sleep, but they do provide context for the rest of the results in the paper. In all samples, leisure and sleep are substitutes, and in all but one case, this is the strongest relationship with sleep. Effects on work time are smaller in magnitude than those on sleep time in both the short and long run, but such changes nonetheless could influence our results using usual weekly earnings as a dependent variable. In sections VB and VD, we obtain similar results when we control for a quadratic in usual hours worked or use only the sample of workers who report an hourly wage, indicating that wage changes are quantitatively more important than changes in work time.

B. *Effect of Sunset on Sleep and Earnings*

Table 3 shows results from estimating equations (11) and (12) on daily ATUS data. The first column shows that the sun setting 1 hour later within a location reduces nighttime sleep by roughly 20 minutes per week, which is statistically significant at the 1% level. We observe about 5 hours of variation in daily sunset time across our sample, meaning that we identify 1.9 hours per week of intra-annual sleep

[12] The sleep definition is the same as that used throughout the rest of the paper. Work time is all time that the individual reports spending at any job (ATUS two-digit time category 05). Home production time is the sum of time spent in personal care (category 01); housework (category 02); caring for a household or nonhousehold member (categories 03 and 04); in education (category 06); shopping (category 07); or using professional, household, or government services (categories 08, 09, and 10). Leisure time is all remaining time categories (11 through 18), including eating, socializing, sports and recreation, attending religious activities, volunteering, and traveling.

[13] In appendix section 3.1, we theoretically investigate reverse causality flowing from wages to sleep, showing that if sleep is relatively more work productive than home productive, an exogenous increase in wages can lead to a decrease in sleep. Consistent with the model, tables 1 and 2 show that short-run sleep is relatively more work productive, and an OLS regression of wages on sleep for this sample yields a negative coefficient. In contrast, sleep appears to be relatively more home productive in the long run, and an OLS regression of wages on average sleep in that sample yields a positive coefficient.

TABLE 3.—SHORT-RUN EFFECTS OF SUNSET ON SLEEP, EARNINGS, AND WAGES FROM ATUS

	(1) First-Stage Sleep	(2) Reduced-Form ln(earnings)	(3) First-Stage Sleep	(4) Reduced-Form ln(wage)
Daily sunset	−0.38***	−0.0044***	−0.42***	−0.0020
	(0.041)	(0.0017)	(0.061)	(0.0016)
Individual controls	Yes	Yes	Yes	Yes
Time controls	Yes	Yes	Yes	Yes
Location FEs	Yes	Yes	Yes	Yes
Mean of dependent variable	57.9	6.67	58.5	2.68
Sample	Earnings	Earnings	Hourly wage	Hourly wage
Observations	61,161	61,161	32,040	32,040
Adjusted R^2	0.12	0.30	0.10	0.35

The table shows results from estimating equation (11), first column, and equation (12), second column, on ATUS data. The dependent variable is indicated at the top of each column. "Earnings" refers to "usual weekly earnings." Sleep is measured in hours per week and sunset time in hours. Controls are discussed in section IIIA and are location fixed effects; race indicators; age indicators; a gender indicator; indicators for holiday, day of week, and year; and occupation indicators. Standard errors, clustered at the FIPS code (location) level, are reported in parentheses. Significance indicated by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.1$.

variation. Practically, this represents a substantial change in time use—roughly equal in magnitude to the 2.1 weekly work hours lost by the average individual during the most recent recession (Aguiar, Hurst, & Karabarbounis, 2013).[14] The time use results in tables 1 and 2 show that this short-run effect is uniform across employed and nonemployed individuals, and unreported heterogeneity analyses show that worker sleep responds fairly equally to solar cues regardless of gender, age, race, or other demographic characteristics. In the appendix, we compare sleep distributions in early- and late-sunset diaries. Early seasonal sunsets shift the whole distribution rightward, increasing sleep for most quantiles. The second column of table 3 shows that daily sunset time also affects earnings in a location. A sunset time 1 hour later reduces earnings by a statistically significant 0.44%.

For workers who report an hourly wage, the first-stage estimate is nearly identical to the full sample estimate. The reduced-form effect of sunset time on wages is less than half the size of the effect on earnings, consistent with the rigidity-attenuation model discussed in section IIIB. In the appendix, we derive exact attenuation predictions based on wage and earnings flexibility estimates from Barattieri, Basu, and Gottschalk (2014). Based on the analysis, we expect the nonhourly worker estimate to be 1.4 times the size of the hourly worker estimate. Estimates comparing earnings effects for the nonoverlapping samples of workers who report an hourly wage and workers who do not, reported in the appendix, show that the nonhourly worker reduced-form estimate is 1.5 times larger. We find additional corroborating evidence for the attenuation model when we examine the reduced-form estimates for union members and government employees. Both of these workers have higher wage rigidity than the typical worker, so our measurement error model predicts that the reduced-form estimate of sunset on earnings will be attenuated for these workers. Indeed, estimates in appendix table 3 show that this is the case.

[14] Our results suggest that the sunset time influences sleep across the sample. Recession-induced work time changes are concentrated among those who lose their jobs, so individual responses to these two shocks will likely differ even though the averages are the same.

Both estimates include individual race indicators (white, black, Asian, and other), age, age squared, a gender indicator, and detailed occupation code indicators (23 categories). Time controls are an indicator for holidays, separate indicators for each day of the week, and year fixed effects. The location fixed effects are at the most disaggregated FIPS code level available for each observation (county, CBSA, or state). The fixed effects absorb any spatial differences in sunset time, leaving only the seasonal component with which to identify the coefficient of interest.

We cluster standard errors at the FIPS code level, as the exogenous variation is at the group rather than the individual level. As shown in the robustness checks below, clustering at higher levels does not change the inference. Other robustness checks related to concerns about seasonal confounders are in section VC.

Table 4 presents estimates of the long-run effects of sunset on sleep (equation [13]), earnings, and wages (equation [14]) using a cross-section of continental U.S. locations from ATUS. In the United States, time zones create just over an hour of variation in long-run sunset time across locations. Thus, average sunset induces about 1 hour per week of sleep variation—roughly half the variation induced by the short-run instrument. Column 1 shows that average weekly sleep falls by just under 1 hour in a location where the sun sets 1 hour later. This is similar to the result obtained by Giuntella, Han, and Mazzonna (2015) using Chinese data. In the appendix, we obtain a roughly similar first-stage estimate using U.S. data collected by Jawbone, a manufacturer of wearable health and sleep trackers. We also show that the sleep difference arises because later sunset time delays sleep onset by more than it delays awakening. These Jawbone estimates indicate that our ATUS results do not arise from reporting biases in the time diaries. In the appendix, we compare sleep distributions in early- and late-sunset locations. Early long-run average sunsets move mass out of the lower tail of sleep and into the center of the distribution. The second column of table 4 shows that for a location where average sunset is 1 hour later, earnings are more than 4% lower on average. To evaluate the relative importance of wage and hours effects, we also estimate the models

TABLE 4.—LONG-RUN EFFECTS OF SUNSET ON SLEEP, EARNINGS, AND WAGES FROM ATUS

	(1) First-Stage Sleep	(2) Reduced-Form ln(earnings)	(3) First-Stage Sleep	(4) Reduced-Form ln(wage)
Average sunset time	−0.93***	−0.045***	−0.73*	−0.061***
	(0.28)	(0.017)	(0.40)	(0.020)
Geographic controls	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes
Mean of dependent variable	57.9	6.67	58.5	2.68
Sample	Earnings	Earnings	Hourly wage	Hourly wage
Observations	529	529	529	529
Adjusted R^2	0.125	0.811	0.097	0.565

The table shows results from estimating equation (13), columns 1 and 3, and equation (14), columns 2 and 4, on ATUS data. The dependent variable is indicated at the top of each column. "Earnings" refers to "usual weekly earnings" and wage to hourly wage for those workers who reports one. Sleep is measured in hours per week and sunset time in hours. Controls are discussed in section IIIA and are an indicator for coastal county, coastal distance, and their interaction; a ten-piece linear spline in latitude; mean age and mean squared age; percent female; race and occupation shares; and a five-piece linear spline in population density. White heteroskedasticity-robust standard errors are reported in parentheses. Significance indicated by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.1$.

for workers who report an hourly wage. Column 3 reports the first stage for wage workers, and column 4 reports the reduced form. Precision is reduced, in part because there are fewer individual observations underlying the location-level averages. Point estimates are similar to those from earnings models, and we cannot reject null hypotheses of equal coefficients across wage and earnings models. While the sample of hourly wage workers is different from the full sample, this pattern of results provides some evidence that wage changes are driving most of the estimated effects on earnings.

We report heteroskedasticity-robust standard errors (White, 1980). We do not cluster because annual average sunset varies exogenously across locations; however, we show in section VC that our results are robust to clustering standard errors at the state level. Section VC also includes a more focused discussion of robustness to potential spatial confounders. While weighted OLS is plausibly more efficient in this setting, unweighted results are presented and discussed in the appendix.

Together tables 3 and 4 demonstrate that sunset time has important effects on time allocation and labor markets. These results are novel and do not require an assumed exclusion restriction. Comparing the two reduced-form estimates, one sees the short-run estimate is about 10% the size of the long-run estimate. While these estimates may differ for many reasons, including different complier populations, this is within the 0 to 25% range suggested by our measurement error analysis, as discussed in section IIIB.

The cross-sectional results pool locations from all four continental U.S. time zones. If our hypothesized causal relationships between sunset, sleep, and wages hold, however, we would expect to find similar results within each time zone. The incomplete geographic coverage of ATUS limits our ability to explore within time zones, so we turn instead to the QCEW, which includes all U.S. counties. The added spatial richness of QCEW also allows for semiparametric estimates of the reduced-form relationship between average sunset time and wages, shown in figure 3. The figure shows a locally weighted kernel regression of residual log earnings on residual average sunset time within each time zone.

To arrive at the figure, we first residualize log wage and average sunset time using a control set similar to our preferred ATUS cross-sectional specification. Controls are time zone indicators, interactions of those indicators with linear splines in coastal distance for the Pacific and Eastern time zones, a linear spline in latitude, median age, percent female, percent black, and percent white. We plot a separate kernel regression of the relationship between this residualized wage and residualized average sunset time for each time zone. There are some striking nonlinearities at the edges of the plot, where data are sparse. In the central region, however, all four time zones exhibit similar negative, roughly linear relationships.[15] In section VB, we present the pooled parametric analog of these QCEW analyses, together with robustness checks.

FIGURE 3.—LONG-RUN EFFECTS OF SUNSET ON WAGES IN QCEW

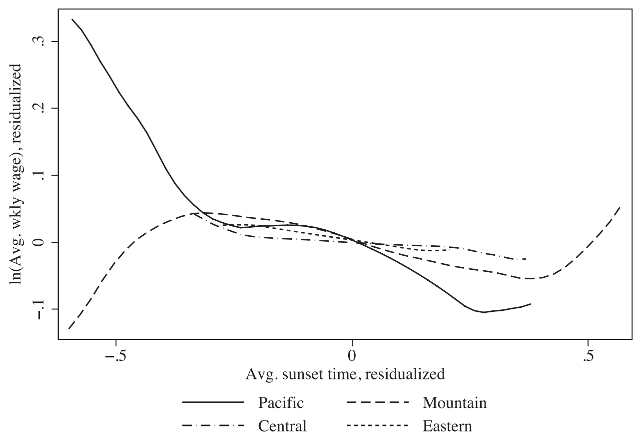

Underlying wage data are a cross-section in locations from QCEW, 1990–2013. The figure shows the relationship between residualized log average wage and residualized sunset time for counties in the Eastern, Central, Mountain, and Pacific time zones. Time zone borders, defined as longitudes at which multiple time zones overlap, are excluded for the reasons discussed in section III7. Residualization is with respect to coastal distance, latitude, median age, percent female, percent black, and percent white. Demographic controls are from 2010 Census data.

C. Robustness

The following sections explore the robustness of our reduced-form specifications. In general, we cannot reject

[15] This suggests that differential attenuation from the larger counties in the Pacific and Mountain zones is not practically important. The relationship between within-time zone longitude and wage is shown in the appendix.

Table 5.—Robustness of ATUS Short-Run Reduced-Form Estimate

	ln(earnings)			ln(earnings)	
Temperature control			*Quarter FEs*		
Daily sunset	−0.0049**	(0.0025)	Daily sunset	−0.0086**	(0.0034)
Temperature and rain			*No holiday season*		
Daily sunset	−0.0048*	(0.0025)	Daily sunset	−0.0049**	(0.0020)
			Observations	52,856	

The table shows results from estimating equation (12) estimated on ATUS data. The dependent variable is indicated at the top of each column. Unless otherwise noted, controls, number of observations, and standard errors are the same as in table 3. Significance indicated by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.1$.

a null hypothesis of equality with our preferred estimates. The differences in coefficients are sometimes large in proportional terms but uniformly small in absolute terms. In general, we can reject a null hypothesis that the coefficient on sunset time is zero, and we comment on the exceptional cases.

Short-run robustness. The primary identification concern with the short-run estimates is that another seasonal variable might be omitted from the sunset time-earnings regressions. We can address this concern in four ways. First, we introduce another seasonal variable, daily temperature, that is plausibly correlated with earnings and covaries with sunset time. We merge temperature data from the NCEP/NCAR reanalysis produced by Kalnay et al. (1996) at the day-location level. The data are available on a two-by-two degree latitude-longitude grid, and we use the daily average temperature for the grid point corresponding to our geocoded centroid for estimation.

The reduced-form coefficients including weather controls are the first two entries in table 5. All of the baseline controls are also included in these regressions. Including temperature changes the coefficient estimate by less than one-third of a standard error. While the temperature control reduces precision, the estimate remains statistically significant at the 5% level. Temperature is highly correlated with sunset time, and it has a small, nominally significant effect on earnings. Adding rainfall reduces significance slightly but does not change the point estimate, since rain is not strongly correlated with either sunset time or earnings. We have also investigated alternative weather controls, including diurnal temperatures and wind speed. Daily average temperature and minimum or maximum temperatures are highly correlated, so the effect on inference is similar. Wind speed does not change the result.

Next, we can include seasonal fixed effects or other calendar controls. The reduced-form estimate including quarter fixed effects is shown in the third entry in the table. The estimate is larger, although precision again declines. With month fixed effects, the first stage becomes 0 to the second decimal place, and the reduced form also becomes 0 to the fourth decimal place. Within a month, there is not enough variation to identify either of these regression coefficients. The fourth entry in table 5 removes the holiday season (Thanksgiving through January 15), again with little change. In the appendix, we show that dropping the entire first and fourth quarters does not substantially change inference; the results show that the reduced-form effect is strongest in quarters 2 and 3.[16]

We can also parsimoniously assess seasonal confounding by including control variables that have the same seasonal pattern as sunset time but different phase or frequency. Figure 2a shows that sunset time follows a sinosoidal pattern within a year, peaking around June 21 and reaching a trough around December 21. The appendix contains results from including seasonal control variables with this same pattern but with peaks and troughs on other days of the year. The results show that the sign of the baseline estimate is robust to including any phase-shifted seasonal control variable of this type.

We find no evidence of confounding seasonal changes in ATUS sample composition. Appendix figure 4a shows that occupation shares in our sample are constant over the months of the year, and appendix figure 4b shows that the share of ATUS respondents reporting a positive wage (i.e., the fraction engaged in market work) is likewise constant over the year. Together, these figures reassure us that such selection bias is not driving our results. Additional robustness checks are reported in appendix section 4.4.

Long-run robustness. We now test the sensitivity of our long-run reduced-form estimate. Again, the first-stage estimates are stable across specifications and are omitted. Table 6 and appendix table 15 show estimates of the reduced-form equation (14) with variations in controls, sample, and clustering. First, a linear control for longitude does not change the estimate substantially. The next robustness check mimics the specification estimated using QCEW data and reported in figure 3. We drop counties at border longitudes where time zones overlap, as such counties might have selected into a time zone based on economic considerations. The standard error is 41% larger, so the test is inconclusive. The point estimate is not statistically distinguishable from 0, or from our preferred estimate. Adding time zone indicators yields an estimate that is statistically significant at the 10% level and similar to our main result. Clustering standard errors at the state level for all observations does not change the results of our hypothesis tests.

[16] In unreported results, the estimate is robust to dropping any two quarters from the sample. Dropping three quarters removes too many data for precise inference, but the point estimates are similar.

Table 6.—Robustness of ATUS Long-Run Estimates

	ln(earnings)			ln(earnings)	
Longitude control			*Albouy QOL control*		
Average sunset time	−0.042**	(0.018)	Average sunset time	−0.032*	(0.019)
No time zone border counties			*No Eastern time zone*		
Average sunset time	−0.037	(0.024)	Average sunset time	−0.049*	(0.029)
Observations	450		Observations	244	
Time zone indicators			*No high-wage cities*		
Average sunset time	−0.043*	(0.023)	Average sunset time	−0.033*	(0.019)
State clustering			Observations	474	
Average sunset time	−0.045***	(0.015)			

The table shows results from estimating equation (14) estimated on ATUS data. The dependent variable is the log of average earnings. Unless otherwise noted, controls, number of observations, and standard errors are the same as in table 4. Results reported under "No high-wage cities" exclude workers in San Francisco, Los Angeles, Chicago, Boston, and New York. Significance indicated by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.1$.

To test whether long-run average sunset time covaries with local amenities, we include a control for the state-level quality-of-life index from Albouy (2008). Again, the change in the coefficient is less than 1 standard error, and it remains significant at the 10% level. Next, we directly test for spatial confounding by excluding the wealthier, denser Eastern time zone, then excluding selected high-wage cities. Both models return estimates (statistically significant at the 10% level) similar to our baseline model, reinforcing the result from the QCEW sample that this effect is not restricted to urban areas or a single time zone.

We next use QCEW data to test the external validity of the long-run ATUS reduced-form estimate. Table 18 in the appendix presents the linear counterpart to the semiparametric analysis in figure 3, pooling across time zones. The base specification seeks to match our long-run ATUS specification as closely as possible, though directly analogous variables are not available in all cases. The point estimate resulting from these controls, −.056, is similar to the long-run ATUS estimate, −.045. We also check the robustness of the QCEW results in appendix table 18. The pattern of robustness checks is similar to those in table 6. In general the QCEW estimates are similar to those from ATUS, however, and in none of the robustness checks can we reject a null hypothesis of equality with the base estimate.

D. *IV Estimates of the Effect of Sleep on Earnings and Wages*

In the preceding sections, we have discussed potential threats to instrument validity (section IIIB) and conducted empirical tests for violations of the exclusion restriction (section VC). Table 7 presents our instrumental variables estimates of the effect of sleep on earnings. A 1-hour increase in short-run mean weekly sleep in a location increases earnings by 1.1%. A 1-hour increase in long-run mean weekly sleep in a location increases earnings by 4.9%. Together these results demonstrate that exogenous location-level sleep changes have important earnings effects.[17]

In the short-run earnings model, the first-stage F-statistic of 85.9 exceeds the relevant Stock-Yogo critical value of 16.38, so we reject the null hypothesis of weak instruments,

Table 7.—IV Estimates of the Effect of Sleep on Earnings and Wages

	(1) Short-Run ln(earnings)	(2) Long-Run ln(earnings)	(3) Short-Run ln(wage)	(4) Long-Run ln(wage)
Sleep	0.011**	0.049**	0.0047	0.083*
	(0.0048)	(0.024)	(0.0039)	(0.050)
Short-run controls	Yes	No	Yes	No
Long-run controls	No	Yes	No	Yes
Observations	61,161	529	32,040	529
Adjusted R^2	0.22	0.70	0.32	0.14
F-statistic on IV	85.90	11.03	47.41	3.33
Elasticity	0.72	2.58	0.29	2.54

The table shows instrumental variables estimates using ATUS data and based on the first-stage and reduced-form equations (11) and (12) (column 1) and equations (13) and (14) (column 2). The dependent variable is the natural log of weekly earnings or wages, as indicated. Controls for columns 1 and 3 are the same as in table 3 and for column 2 and 4 are the same as in table 4. Significance indicated by $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.1$.

where "weak" is defined as true size greater than 10% for a nominal 5% test (Stock & Yogo, 2005).[18] In the long-run earnings model, the first-stage F-statistic of 11.03 falls short of the Stock-Yogo critical value for true size of 10% but exceeds the critical value of 8.96 for true size of 15%. The result of our long-run t-test should thus be viewed with more skepticism, but we can rule out gross failures of size control.

While our earnings estimates largely reflect wage changes rather than changes in hours worked, table 7, columns 3 and 4, directly examines effects on hourly wage for workers who report one. A 1-hour increase in short-run mean weekly sleep in a location increases wages by 0.5%. A 1-hour increase in long-run mean weekly sleep in a location increases wages by 8.3%. Precision is reduced relative to the earnings models. The short-run estimate is not statistically significant, and the long-run estimate is statistically significant only at the 10% level. In the long-run wage model, the first-stage F-statistic is quite small at 3.33, and we strongly caution against placing too much weight on this particular point estimate. Consistent with the analysis in appendix section 3.2 and the earnings flexibility estimates of Barattieri et al. (2014), the short-run estimates are smaller relative to the long-run estimates when using wages rather than earnings.

Our estimates reflect location-level local average treatment effects, so several nuances of interpretation warrant

[17] The appendix shows an investigation of possible nonlinearity in the sleep-earnings relationship and finds no evidence for it.

[18] These critical values assume homoskedasticity, so test results are merely suggestive.

discussion. First, all workers in a location experience the same sunset time. If sunset-induced sleep differences generate productivity spillovers across workers, our estimated β will capture not the effect of increasing individual sleep but rather the effect of increasing mean sleep in a location. The empirical literature on peer effects provides evidence that such spillovers may be large. For example, Chetty et al. (2011) estimate an elasticity of age 27 earnings with respect to kindergarten peer quality of roughly 1. Similarly, Carrell, Fullerton, and West (2009) estimate an elasticity of academic achievement with respect to peer quality of roughly 0.9. Mas and Moretti (2009) find evidence of large, positive peer effects in supermarket cashiers. Moretti (2004) finds that human capital spillovers can operate at the municipal level, estimating that a 1% increase in the share of college-educated workers in a city increases output by roughly half of 1%. This is similar in spatial scale to our analysis. It suggests that if earlier sunset time increases sleep and productivity for all workers in a location, spillover effects may be broadly felt. Because spillovers and peer effects are externalities, it would not necessarily be individually rational for a worker to sleep more. This body of evidence implies that an individual worker would not see a 4.9% wage increase from an additional hour of sleep, but instead something potentially much smaller. The second interpretative nuance is that differences in long-run sunset time across locations are nearly permanent. For most locations, long-run sunset time has changed only when the United States has revised its time zones, which it has not done since World War I. If labor and capital are on average complements, it is possible that sleep-driven labor productivity differences have influenced the long-run growth of the capital stock. Thus, the effects we estimate may have emerged over many years, and short-run effects could be different. Third, a positive marginal effect does not imply optimization failure. If the marginal willingness to pay for leisure exceeds the marginal wage increase from additional sleep, an optimizing worker will not increase sleep.

Our analysis demonstrates that workers experiencing an earlier sunset get more sleep. As discussed in section VA, in the short run, the additional sleep largely comes at the expense of leisure, while in the long run, it comes at the expense of both work and leisure. Insofar as these changes in other time uses affect worker productivity, our instrumental variables estimates of the effect of sleep on wages will also contain those effects. In other words, we estimate the effect of a quasi-random change in sleep, allowing other time uses to adjust endogenously. While this might seem undesirable at first glance, it is unavoidable. An agent's time constraint always binds with perfect equality. Even in a laboratory setting, it is not possible to change the time use of interest without also changing at least one other time use.

Expressed as an elasticity, our short-run earnings estimate is 0.84, and our long-run earnings estimate is 2.6. Medical researchers have typically found elasticities of task performance with respect to sleep duration of approximately four (table 1). If wages are equal to a worker's marginal physical product multiplied by output price, we expect such performance effects to produce equally large wage effects. Our smaller estimated elasticities may reflect differences between laboratory tasks and actual work tasks or the broader scope for adaptation (e.g., the use of stimulants like coffee) outside the lab. The difference between short- and long-run elasticities is consistent with the attenuation bounds calculated in the appendix. Unaided intuition might suggest smaller effects of sleep on performance, but intuition provides a poor sense of this relationship. Van Dongen et al. (2003) showed that subjects' self-reported fatigue quickly stabilized after a few days of sleep reduction, even as their performance continued to decline.

Taking average values for earnings and assuming fifty workweeks per year, one can calculate the annual income effects implied by our long-run estimates. If mean weekly sleep in a location increased by 1 hour and work time remained unchanged, mean annual income would rise by about $2,350. In reality, extra sleep comes out of both work and nonwork time. If workers took roughly 70% of the extra sleep hour out of work time, then a 1-hour increase in weekly mean sleep in a location would increase mean annual income by about $1,570. If extra sleep came solely at the expense of work time, the income increase would be $1,250. As we showed analytically in section IIB, an optimizing agent might in fact choose to increase work time in response to an increase in sleep.

VI. Conclusion

Although time use is entangled in a causal web with labor market outcomes, economists have paid little attention to these relationships. In particular, the profession has scarcely examined sleep. In this paper, we formalize the links between productivity-improving sleep, wages, and other time uses in an optimizing model. We generate theoretical predictions on complementarity between sleep and other time uses for both workers and nonworkers. Empirical tests yield results largely consistent with those predictions. Our earnings and wage regressions demonstrate that sleep has a powerful impact on labor market outcomes and should be considered an integral part of a worker's utility maximization problem. Using individual time use diaries matched with labor market variables from ATUS, we show that increasing short-run weekly average sleep in a location by 1 hour increases worker earnings by roughly 1%. Increasing long-run weekly average sleep in a location by 1 hour increases earnings by roughly 5%. Our use of instrumental variables techniques addresses the reverse causality and omitted variable problems that would bias naive estimates. We buttress this finding with a battery of short- and long-run robustness checks and a hedonic model of home prices showing that long-run wage increases are partially capitalized into housing.

These results suggest that sleep is a crucial determinant of wages, rivaling ability and education in importance. Figures of the magnitude shown here naturally lead one to ask why workers do not work less and sleep more. One possible explanation lies in the spillovers and general equilibrium effects our estimate incorporates. It is also conceivable that observed sleep reflects optimization failure by workers, as hypothesized by Mullainathan (2014). Such failure could occur even under classical assumptions. For example, the inaccurate self-perceptions of fatigue found by Van Dongen et al. (2003) could lead to sleep below the utility-optimizing level even if workers are behaving optimally, conditional on their information set. Such suboptimal sleep could contribute to the type of time use poverty trap analyzed by Banerjee and Mullainathan (2008).[19] On the other hand, suboptimal sleep could arise from behavioral considerations like time inconsistency or constraints on cognitive resources (see, e.g., Mani et al., 2013). While optimization failure and its possible mechanisms are beyond this scope of this paper, we are exploring them in ongoing experimental work.

Further attention should be paid to industries characterized by chronic sleep shortages. In addition to wages, optimal sleep plausibly depends on other factors like leisure complementarities, direct sleep utility, and health optimization. Each of these trade-offs suggests an interesting research question. More broadly, our results demonstrate that nonlabor time uses can have first-order effects on labor outcomes, and these effects warrant further investigation.

[19] While we interpret their work as a model of inattention, they note, "In fact, our model is formally identical to a rational time allocation model, if we think of comfort goods as time saving devices."

REFERENCES

Aguiar, Mark, and Erik Hurst, "Measuring Trends in Leisure: The Allocation of Time over Five Decades," *Quarterly Journal of Economics* 122 (2007), 969–1006.
Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis, "Time Use during the Great Recession," *American Economic Review* 103 (2013), 1664–1696.
Albouy, David, "Are Big Cities Bad Places to Live? Estimating Quality of Life across Metropolitan Areas," NBER working paper 14472 (2008).
Allen, Roy, and John Rehbeck, "Hicksian Complementarity and Perturbed Utility Models" (2016). https://ssrn.com/abstract=280135.
Allen, W. F. "Open Letters: Standard Railway Time," *Century* (1883), 796.
Banerjee, Abhijit V., and Sendhil Mullainathan, "Limited Attention and Income Distribution," *American Economic Review* 98 (2008), 489–493.
Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk, "Some Evidence on the Importance of Sticky Wages," *American Economic Journal: Macroeconomics* 6 (2014), 70–101.
Barnes, Christopher M., and David T. Wagner, "Changing to Daylight Saving Time Cuts into Sleep and Increases Workplace Injuries," *Journal of Applied Psychology* 94 (2009), 1305.
Becker, Gary S., "Investment in Human Capital: A Theoretical Analysis," *Journal of Political Economy* 70 (1962), 9–49.
——— "A Theory of the Allocation of Time," *Economic Journal* 75 (1965), 493–517.
——— *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education* (Chicago: University of Chicago Press, 1964).
Biddle, Jeff E., and Daniel S. Hamermesh, "Sleep and the Allocation of Time," *Journal of Political Economy* 98 (1990), 922–943.
Black, Sandra E., and Lisa M. Lynch, "How to Compete: The Impact of Workplace Practices and Information Technology on Productivity," this REVIEW 83 (2001), 434–445.
Bloom, Nicholas, Raffaella Sadun, and John Van Reenen, "Americans Do It Better: US Multinationals and the Productivity Miracle," *American Economic Review* 102 (2012), 167–201.
Bonke, Jens, "Do Morning-Type People Earn More Than Evening-Type People? How Chronotypes Influence Income," *Annales d'Economie et de Statistique* 105–106 (2012), 55–76.
Brochu, Pierre, Catherine Deri Armstrong, and Louis-Philippe Morin, "The 'Trendiness' of Sleep: An Empirical Investigation into the Cyclical Nature of Sleep Time," *Empirical Economics* 43 (2012), 891–913.
Bureau of Labor Statistics, "Employment and Wages Covered by Unemployment Insurance," in *BLS Handbook of Methods* (Washington, DC: Bureau of Labor Statistics, 1997).
——— *American Time Use Survey User's Guide* (2015).
Cappuccio Francesco, P., Lanfranco D'Elia, Pasquale Strazzullo, and Michelle A. Miller, "Sleep Duration and All-Cause Mortality: A Systematic Review and Meta-Analysis of Prospective Studies," *Sleep* 33 (2010), 585–592.
Carrell, Scott E., Richard L. Fullerton, and James E. West, "Does Your Cohort Matter? Measuring Peer Effects in College Achievement," *Journal of Labor Economics* 27 (2009), 439–464.
Chetty, Raj, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane W. Schanzenbach, and Danny Yagan, "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR," *Quarterly Journal of Economics* 126 (2011), 1593–1660.
Frazis, Harley, and Jay Stewart, "How to Think about Time-Use Data: What Inferences Can We Make about Long- and Short-Run Time Use from Time Diaries?" *Annals of Economics and Statistics* 105–106 (2012), 231–245.
Giuntella, Osea, Wei Han, and Fabrizio Mazzonna, "Circadian Rhythms, Sleep and Cognitive Skills: Evidence from an Unsleeping Giant," IZA working paper 9774 (2015).
Graff Zivin, Joshua and Matthew Neidell, "The Impact of Pollution on Worker Productivity," *American Economic Review* 102 (2012), 3652–3673.
Gronau, Reuben, "Leisure, Home Production, and Work: The Theory of the Allocation of Time Revisited," *Journal of Political Economy* 85 (1977), 1099–1123.
Hamermesh, Daniel S., Harley Frazis, and Jay Stewart, "The American Time-Use Survey," *Journal of Economic Perspectives* 19 (2005), 221–232.
Hamermesh, Daniel S., Caitlin Knowles Myers, and Mark L. Pocock, "Cues for Timing and Coordination: Latitude, Letterman, and Longitude," *Journal of Labor Economics* 26 (2008), 223–246.
Hubert, Marc, Marie Dumont, and Jean Paquet, "Seasonal and Diurnal Patterns of Human Illumination under Natural Conditions," *Chronobiology International* 15 (1998), 59–70.
Ian Bartky. "The Adoption of Standard Time," *Technology and Culture* 30 (1989), 25–56.
Kalnay, Eugenia, Masao Kanamitsu, Robert Kistler, William Collins, Dennis Deaven, Lev Gandin, Mark Iredell, Suranjana Saha, Glenn White, John Woollen, et al. "The NCEP/NCAR 40-Year Reanalysis Project," *Bulletin of the American Meteorological Society* 77 (1996), 437–471.
Lambert, G. W., C. Reid, D. M. Kaye, G. L. Jennings, and M. D. Esler, "Effect of Sunlight and Season on Serotonin Turnover in the Brain," *Lancet* 360 (2002), 1840–1842.
Leibenstein, Harvey, *Economic Backwardness and Economic Growth* (New York: Wiley, 1957).
Library of Congress. "Time!" November 2010. http://memory.loc.gov/ammem/today/nov18.html.
Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao, "Poverty Impedes Cognitive Function," *Science* 341:6149 (2013), 976–980.
Mas, Alexandre, and Enrico Moretti, "Peers at Work," *American Economic Review* 99 (2009), 112–145.
Meeus, Jean H., *Astronomical Algorithms* (Richmond, VA: Willmann-Bell, 1991).
Moretti, Enrico, "Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions," *American Economic Review* 94 (2004), 656–690.
Mullainathan, Sendhil, "Get Some Sleep, and Wake Up the G.D.P.," *New York Times*, February 1, 2014.

Murase, Sanako, M. Kitabatake, T. Yamauchi, and A. A. Mathe, "Seasonal Mood Variation among Japanese Residents of Stockholm," *Acta Psychiatrica Scandinavica* 92 (1995), 51–55.

Mushkin, Selma J. "Health as an Investment," *Journal of Political Economy* 70 (1962), 129–157.

"The New Standard Time and Its Advantages," *Manufacturer and Builder* 15 (1883).

Psacharopoulos, George, and Harry Anthony Patrinos, "Returns to Investment in Education: A Further Update," *Education Economics* 2004 (2), 111–134.

Roenneberg, Till, C. Jairaj Kumar, and Martha Merrow, "The Human Circadian Clock Entrains to Sun Time," *Current Biology* 17 (2007), R44–R45.

Schober Michael F., and Frederick G. Conrad, "Does Conversational Interviewing Reduce Survey Measurement Error?" *Public Opinion Quarterly* 61 (1997), 576–602.

Smith Austin, C., "Spring Forward at Your Own Risk: Daylight Saving Time and Fatal Vehicle Crashes," *American Economic Journal: Applied Economics* 8 (2016), 65–91.

Solon, Gary Steven, J. Haider, and Jeffrey M. Wooldridge, "What Are We Weighting For?" *Journal of Human Resources* 50 (2015), 301–316.

Stock, James H., and Motohiro Yogo, "Testing for Weak Instruments in Linear IV Regression" (p. 80), in Donald W. K. Andrews and James H. Stock, eds., *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg* (Cambridge: Cambridge University Press, 2005).

Szalontai, Gábor, "The Demand for Sleep: A South African Study," *Economic Modelling* 23 (2006), 854–874.

Taras, Howard, and William Potts-Datema, "Sleep and Student Performance at School," *Journal of School Health* 75 (2005), 248–254.

Ulmer, Cheryl, Dianne Miller Wolman, and Michael E. Johns, eds., *Resident Duty Hours: Enhancing Sleep, Supervision, and Safety* (Washington, DC: National Academies Press, 2009).

Valpando Adonna, "DOT Procedure for Moving an Area from One Time Zone to Another" (2013). https://www.transportation.gov/regulations/procedure-moving-area-one-time-zone-another.

Van Dongen, P. A. Hans, Greg Maislin, Janet M. Mullington, and David F. Dinges, "The Cumulative Cost of Additional Wakefulness: Dose-Response Effects on Neurobehavioral Functions and Sleep Physiology from Chronic Sleep Restriction and Total Sleep Deprivation," *Sleep* 26 (2003), 117–126.

White, Halbert, "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica* 48 (1980), 817–838.

Winship, Christopher, and Larry Radbill, "Sampling Weights and Regression Analysis," *Sociological Methods and Research* 23 (1994), 230–257.