Work Effort and the Cycle: Evidence from Survey Data*

Vivien Lewis
Deutsche Bundesbank

David Van Dijcke
University of Oxford

October 9, 2019

Abstract

We use data from the World Values Survey and the Work Orientations Survey to analyse the cyclical nature of work effort and attitudes to work effort. Our aim is to test two competing theories of labor effort, the labor hoarding view and the Shapiro and Stiglitz (1984) 'shirking model'. Self-reported work effort is found to be strongly procyclical, while attitudes to effort move slightly countercyclically. We provide evidence for the presence of labor hoarding by showing how the cyclicality of effort changes with the strictness of employment protection legislation. Finally, we document heterogeneity in effort cyclicality across occupations and individuals.

Keywords: effort, labor hoarding, labor productivity, shirking, World Values Survey, Work Orientations Survey.

JEL classification: E24, E32, E71.

*Contact address: Research Centre, Deutsche Bundesbank, Wilhelm-Epstein-Str. 14, 60431 Frankfurt am Main, Germany. Phone: +49 (0)69 9566 6254. E-Mail: vivien.lewis@bundesbank.de. We thank Olga Goldfayn-Frank, Matthias Hertweck, Valentin Kecht, Elmar Mertens, Fabien Postel-Vinay, Panagiota Tzamourani, Harald Uhlig, Michèle Tertilt, Tom Holden, as well as seminar participants at the Bundesbank for helpful discussions and comments. The views expressed in this paper are those of the authors and do not necessarily coincide with the views of the Bundesbank or the Eurosystem.

1 Introduction

What drives the cyclical component of effort at work? In this paper, we look for answers to this question in two surveys: the World Values Survey and the Work Orientations Survey. Two models of labor effort are particularly prominent in the economics literature: the 'shirking model' of Shapiro and Stiglitz (1984) and the labor hoarding hypothesis first put forth by Solow (1964). In the first model, effort results from the fear of lay-off when caught shirking on the job. Effort is countercyclical; it is high in times of high unemployment when – accordingly – the job finding rate is low. In the second model, effort adjusts to temporary changes in demand when hiring and firing (or overtime hours) are costly. This labor hoarding view thus posits that effort is procyclical.

In many countries and over different time periods, labor productivity is procyclical. The labor hoarding view is consistent with this observation. However, competing explanations exist, such as the prevalence of technology shocks in driving business cycles, intangible capital investment, or variable capital utilization. Thus, procyclical productivity does not necessarily imply procyclical effort. Here, we try to circumvent this problem by measuring self-reported work effort, as well as attitudes to effort, using evidence from the Work Orientations Survey (WOS) and the World Values Survey (WVS), respectively. This gives us an idea of the forces that drive effort and allows us to discriminate between the two theories: shirking vs. labor hoarding.

Finding an answer to our research question is relevant in a number of ways. First, to the extent that effort acts as a margin of labor adjustment in response to demand shocks, it can be used to smooth business cycle fluctuations. In particular, labor-hoarding firms can vary hours per worker and effort per hour when demand for their products changes, attenuating unemployment fluctuations. This is especially useful in countries where employment protection legislation makes hiring and firing costly. As shown by Ohanian and Raffo (2012), such countries indeed rely much more on the intensive labor margin than those with very fluid labor markets like the US. As a consequence, the necessity of macroeconomic stabilization policies is reduced. Additionally, evidence for procyclical labor effort adjustments would bolster the case of unemployment stabilization policies that "subsidize labor hoarding" (Giupponi and Landais, 2018). Second, identifying the sources of procyclical labor productivity has implications for our assessment of the nature of shocks driving the business cycle.[1] Unobserved procyclical labor utilization (or effort) is one possibly explanation. Adherents of the Real Business Cycle (RBC) paradigm tend to attribute rising productivity to technological improvements. This conclusion is no longer inevitable when the utilization of inputs is variable, as demand shocks can in that case be consistent with procyclical productivity as well. Third, evidence for the shirking model of effort supports the theory of efficiency wages, which can help explain why firms pay wages far above their employees' reservation wages.

Our estimation results indicate that attitudes to effort at work move countercyclically, while self-reported effort at work moves procyclically.

The outline of the paper is as follows. Section 2 discusses the two models in greater detail and summarizes the existing empirical evidence on the topic. Section 3 briefly introduces the World Values Survey and the Work Orientations Survey. In Section 4, we explain how we measure (attitudes to) work effort from participants' responses to selected questions. We show correlations at the country-wave level and regression results at the individual level. Section 5 concludes.

[1]See Basu and Fernald (2001) and Fernald and Wang (2016) for overview articles.

2 Shirking vs. Labor Hoarding

Below, we summarize the empirical and theoretical contributions to the literature on the two competing theories of effort, labor hoarding and shirking.

2.1 Empirical Literature

Okun (1963) made the key observation that, while unemployment is negatively related to output, the estimated coefficient is systematically below one (in absolute value). In an expansion, productivity thus rises alongside the fall in unemployment. As is clear from the following quote, Okun believed that variable labor utilization was the force behind the procyclicality of measured labor productivity:

> *The record clearly shows that manhour productivity is depressed by low levels of utilization [...] Indeed, many a priori arguments have been made for the reverse view – that depressed levels of activity will stimulate productivity through pressure on management to cut costs, through a weeding-out of inefficient firms and low quality workers, and through availability of more and higher quality capital per worker for those employees who retain their jobs. If such effects exist, the empirical record demonstrates that they are swamped by other forces working in the opposite direction.*
>
> Okun (1963)

The empirical literature suffers from the problem that labor effort is, usually, unobserved. One might use proxies for effort and examine their cyclical properties. However, this approach requires additional assumptions. For instance, it is conceivable that – due to job-related stress and hazardous working conditions – greater work effort during booms leads to reduced health and more workplace accidents. It has indeed been shown that various relevant indicators are procyclical: mortality rates and unhealthy behavior patterns (Ruhm, 2000), sick leave and absenteeism (see Taylor (1979), Leigh (1985), Schön (2015)), workplace accidents (Fairris (1998), Boone and van Ours (2002)). Moreover, the observation that the procyclicality of mortality has disappeared recently (Ruhm, 2015) happens to coincide with the vanishing procyclicality of labor productivity, see Galí and van Rens (2014). Taken together, this evidence points to work effort being procyclical.

Those rare studies that measure effort directly report conflicting results. Using productivity data from one large (anonymous) US firm, Lazear, Shaw, and Stanton (2016) document a positive correlation between effort and the local unemployment rate. That is, they find evidence for countercyclical effort. Senney and Dunn (2019) present a direct measure of effort from a large automobile manufacturing plant over the 1980s. They find that effort responds positively to a worsening of macroeconomic conditions and conjecture that this effect is driven by the fear of plant closure. Both papers' results pertain to a single firm in the US; it is therefore unclear whether it holds more generally.

The American Time Use Survey (ATUS) instead reveals that effort at work varies procyclically (Burda, Genadek, and Hamermesh, 2019). More precisely, the time spent on 'non-work at work', which is a counterindicator of effort, varies positively with local unemployment. Hence, according to evidence from the ATUS, effort is procyclical. Fay and Medoff (1985) report that during downturns, firms paid for about 8 percent more labor hours than were technologically necessary to meet production requirements. This indicates that part of the workforce exerts low effort during these times. Firms retain workers nonetheless, since the cost of firing, re-hiring and training workers would exceed the costs of idleness.

2.2 Theoretical Literature

In Shapiro and Stiglitz (1984), unemployment acts as a 'worker discipline device': a high unemployment rate implies a reduced chance of finding a job elsewhere if fired. This mechanism induces workers to exert greater effort on the job during recessions. The Shapiro-Stiglitz shirking model assumes that the worker has full discretion over her labor effort. Workers are effort minimizers; only the risk of being caught shirking – and consequently getting the sack – induces them to work rather than shirk. Insofar as the shirking model implies countercyclical labor effort, Uhlig and Xu (1996) argue that the model is incompatible with technology-driven explanations of the business cycle, since it would require implausibly large technology shocks in order to be consistent with the observed procyclicality of labor productivity. Another strand of the theoretical literature views effort as procyclical. Oi (1962) develops the argument that labor is a 'quasi-fixed' factor. High costs of adjusting a firm's workforce lead to the practice of 'labor hoarding'.[2] In the original formulation of Oi (1962), firms only fire a worker if her marginal product of labor falls below her wage. In the long run, this marginal product equals the wage rate plus the amortized sunk training/hiring cost. In the short run, however, the latter costs are sunk, and thus should not enter into the firm's firing decision. The resulting wedge between wage and marginal product explains why in recessions firms will hold on to less productive workers instead of firing them.[3]

Bils and Cho (1994) develop a model with three labor margins, employment, hours per employee and effort per hour worked, which all give rise to disutility. Firms demand the number of effective hours required to produce a given amount of output, and workers can choose the combination of hours spent at the workplace and effort per hour that minimizes disutility. Equilibrium effort is a positive convex function of hours per worker, such that the procyclical hours imply procyclical effort.

Uhlig (2004) proposes a model with 'workplace leisure' and 'home leisure', which also gives rise to procyclical effort, as in downturns home leisure is substituted for workplace leisure. That way, he questions the well-known result of Galí (1999) that technology shocks cannot be the main drivers of business cycles (as in RBC models), since they imply a negative correlation between labor productivity and hours. This negative correlation, Uhlig argues, arises partially from a mismeasurement of labor productivity due to workers spending a share of their work hours on workplace leisure. Nonetheless, several other papers find that introducing variable labor utilization in an RBC setting reduces the importance of technology shocks to the business cycle by diminishing the variance of technology shocks (Burnside, Eichenbaum, and Rebelo, 1993; Basu and Kimball, 1997). In general, then, the existence of both pro- and countercyclical labor effort would seem to put in doubt the claim that technology shocks are the predominant drivers of the business cycle, while demand-driven explanations of the business cycle sit awkwardly with the shirking model of labor effort but align well with the labor hoarding view.

As a side note, the absence of any observed cyclicality in effort must not lead us to

[2]Biddle (2014) provides a review of the emergence of the labor hoarding concept.

[3]In a model without sunk costs, a firm reacts to a drop in output by firing workers, since the decline in the marginal product (unit of output per hour) necessitates a reduction of the wage if the firm's optimality conditions are to remain satisfied.

discard models of endogenous effort. In Danthine and Kurmann (2004), a worker's effort function depends on the compensation she receives; she provides some extra effort if she is paid a 'fair' wage, that is, one that is above some reference level, such as last period's real wage. In equilibrium, the Solow (1979) condition holds; firms find it optimal to set wages so as to elicit a constant effort level. Moreover, Collard and de la Croix (2000) show that the fair wage model may be consistent even with procyclical effort if the reference wage is related to the worker's *own* past wage. Thus, any evidence for procyclical effort should be interpreted as evidence against the shirking model, but not necessarily against efficiency wage theory more generally.

To sum up, the theoretical literature on the cyclicality of effort is just as divided as the empirical evidence. Which of the two views is more relevant: procyclicality, which is consistent with the labor hoarding model, or countercyclicality, which is consistent with the shirking model? In the following, we will tease out an answer to this question from survey data using the WOS and the WVS.

3 Data Description and Methodology

The microeconomic data come from the World Values Survey and the Work Orientations Survey.

3.1 Data Description

Table 1 provides a description of each variable presented in the analysis below, as well as the source it was obtained from.

— [insert Table 1 here] —

In the following, we briefly describe the two main survey data sources used in this study. The macroeconomic data we use is obtained from the OECD, except for hours worked, which are obtained from the Penn World Tables.

World Values Survey. The World Values Survey started in 1981 and has seen six waves to date, the last one having been released in 2014.[4] The survey currently covers almost 100 countries, in each of which it aims to obtain a nationally representative sample of all people between 18 and 85. Sampling methods differ across countries, but are held to rigorous standards, with the minimum allowable sample size being 1,200.[5] The face-to-face interviews are carried out either by means of a paper questionnaire or as a computer-assisted personal interview. The survey aims to minimize non-response, but documents it where it occurs (Inglehart, Haerpfer, Moreno, Welzel, Kizilova, Diez-Medrano, Lagos, Norris, Ponarin, Puranen, et al., 2014).

Since we are mostly interested in OECD countries, we integrate the World Values Survey with the European Values Study, which is maintained by the University of Tilburg and the Data Archive for the Social Sciences in Cologne (EVS, 2015; WVS, 2015). Both surveys use a harmonized dictionary for the data and trend variables, leading to the combined Integrated Values Survey, which currently comprises a total of 364 conducted surveys and more than 500,000 observations. Our particular subsample, which consists

[4]Detailed information on the WVS is provided at www.worldvaluessurvey.org.

[5]'Sample size' here denotes the number of correctly completed interviews.

of all full-time employees in OECD countries, comprises 96,693 observations, with 173 country-years in total.

A downside of the survey is that the sampling methods used in each country are not documented, which precludes the possibility of correcting statistical tests for sample stratification. To address this deficiency, we cluster all our standard errors at the country level. Nonetheless, the reported standard errors should be interpreted with caution, as they are probably biased to the downside.

Papers that have used the WVS to shed light on (macro-)economic questions include Ehrmann and Tzamourani (2012), who consider inflation fears; Doepke and Zilibotti (2017) who look at parenting styles and inequality; Giavazzi, Schiantarelli, and Serafinelli (2013), who estimate the impact of cultural attitudes on employment, and Buch and Engel (2013), who also use the WVS to study labor effort, though they focus on its relationship with preferences for redistribution rather than its cyclical properties.

Work Orientations Survey. The Work Orientations Survey started in 1989 and currently comprises four waves, with later waves always being partial replications of earlier ones. The survey deals with issues related to people's work environment and their subjective experience thereof. It falls under the umbrella of the International Social Survey Programme (ISSP), a collaboration that includes institutional members from various countries. As with the WVS, sampling methods differ across countries, but they are all a form of multi-staged stratified random sampling and this of citizens older than 18 (16 for Japan). Interview methods differ across countries, but are mostly face-to-face. A sample for any given country-wave combination always comprises at least 1,000 data points. In total, our subsample of full-time OECD-based employees is made up of 48,113 observations for 84 country-years.

Note that for Waves II-IV, the sample includes any potential short-time workers. For Wave I, it only includes those full-time workers who usually work at least 30 hours per week. For three reasons, this should not bias our estimates too much (Cahuc and Carcillo, 2011). First, short-time work programs were less wide-spread in 1989 than nowadays; second, the 30h per week cut-off includes at least some short-time workers; third, the number of countries in Wave I is relatively small.

While the WOS does report sampling methods for each country, sampling strata and clusters are not readily available, so standard errors are plagued by the same issue as the WVS.

One paper from the field of public administration that also uses the WOS to proxy labor effort, in the form of the ratio between government and private-sector wages, is Taylor and Taylor (2011).

3.2 Summary Statistics

We are interested in variables that relate to effort at work. An overview of the specific survey questions our various measures are based on can be found in Table 2.

— [insert Table 2 here] —

The use of two separate data sets with various measures of effort at work – both indirectly through attitudes and directly through self-reporting – allows us to approximate labor effort in a rich manner. First, the World Values Survey contains several questions that relate to the respondent's attitude towards her job and towards her work environment more generally. We will refer to the effort measure from the WVS as "attitudes to effort at work". Second, the Work Orientations Survey contains questions that ask directly about the amount of effort the respondent exerts at work, as well as about the concomitant signs of increased exertion, namely stress and exhaustion. For that reason, we will refer to the effort measures obtained from the WOS as "self-reported effort at work". On the basis of these two sets of proxies for labor effort, we are able to provide a detailed picture of its cyclical properties.

In Tables 3 and 4, we provide summary statistics for the main variables used in our analysis, for the WVS and the WOS, respectively.

— [insert Tables 3 and 4 here] —

The main takeaways from these tables for the macro variables are that the WVS sample includes more recessionary country-years than does the WOS, as can be seen from the lower values for the cyclical measures. Also, the output gap has a larger standard deviation than the unemployment rate, a feature that shows up in its regression coefficient.

For the micro variables, a striking and intuitive feature in Table 4 is that high self-reported effort at work, as well as exhaustion and stress due to work, are prevalent. That is, people on average indicate that they work hard and that they are stressed and exhausted more often than "sometimes". When it comes to attitudes to effort, on the other hand, Table 3 shows that respondents indicate effort-related aspects of a job as important slightly more than half of the time, with about half of the respondents indicating that they find opportunity for initiative on the job important. For the control variables, the averages of the marriage rate, education rate and respondents' age are in line with OECD averages. The share of people with only secondary education is quite a bit higher in the WOS than in the WVS, which is due to the WOS including details on those who only finished lower secondary education. The average share of respondents in unions is also in line with the OECD average, although there is a wide divergence between country-specific averages. Female respondents seem to be undersampled in both datasets, but the survey weights should correct for this. In all, the data seems to be representative of the wider population of the OECD countries from which it was sampled.

3.3 Methodology

All regressions are either simple logit models or ordered logit models - so-called "proportional odds models", as first proposed by McCullagh (1980). The latter can be represented as

$$logit(P(Y \leq k|X)) = \theta - X'\beta, \tag{1}$$

where Y is our dependent variable, X is a matrix of explanatory variables and controls, k is the relevant threshold level of the dependent categorical variable, and $logit(P)$ equals the log odds of P, that is $log(\frac{P}{1-P})$. The slope coefficients reported are odds ratios, obtained as e^{β}. This means that coefficients greater than 1 indicate a positive effect of an increase in the explanatory variable in question, x, on the odds ratio

$$\frac{P(Y > k|x)}{P(Y \leq k|x)}, \tag{2}$$

while a coefficient smaller than 1 indicates the opposite. In our context, a coefficient on cyclical unemployment e^{β} greater than 1 thus indicates countercyclical effort (effort increases in a downturn), and vice versa for a coefficient below 1. The standard errors are transformed via the delta method to reflect the uncertainty of the coefficient vis-à-vis the null hypothesis $\beta = 1$. Note that this model estimates only one slope per dependent variable, but $J - 1$ intercepts θ, where J is the number of levels of the dependent categorical variable. In that, the model rests on the so-called proportional odds assumption, which states that

$$\frac{P(Y > k|x)}{P(Y \leq k|x)} - \frac{P(Y > (k+1)|x)}{P(Y \leq (k+1)|x)} \tag{3}$$

is the same for all $k = 1, .., J - 1$. To test this assumption, we report the p-value of the likelihood ratio test for the ordered logit model against the multinomial logit model, since the former can be shown to be nested in the latter. A p-value smaller than 0.05 means that we reject the proportional odds assumption for the model as a whole. In general, the likelihood ratio test does reject the proportional odds assumption for model as a whole. There is, however, no clear reason why the full model should satisfy the assumption, given the presence of fixed effects dummies. Therefore, we also report the p-value of the Brant test for the coefficient on unemployment only, which is of primary interest. The Brant test compares the coefficients of $k - 1$ dichotomized logit models, which transform the dependent variable into a dummy variable that is 1 if $y <= k$ and 0 otherwise, with the coefficient of the ordered logit model, where k is the total number of values the dependent variable (self-reported effort) can take (Brant, 1990, §3). The test rejects when these coefficients are significantly different. In our case, the Brant test generally accepts the proportional odds assumption for the coefficients on unemployment and the alternative cyclical indicators we consider in Table 6. In general then, the proportional odds assumption seems reasonable. We now present the results in detail in the next section.

4 Results

First, we present results at the country-wave level. Second, we dig deeper and conduct individual-level regressions to examine in greater detail the determinants of work effort.

4.1 Country-Level Evidence

As a first pass at examining the data, we conduct a similar type of analysis as do Doepke and Zilibotti (2017). Country averages per survey wave are taken as single observations. Averages are survey-weighted. As explained above, once we have a measure of effort, we can consider its cyclicality in order to discriminate between the two theories. A positive correlation between macroeconomic conditions and work effort would be consistent with the labor hoarding view. A negative correlation between macroeconomic conditions and work effort would instead support the shirking model.

Figure 1 plots the three self-reported effort measures from the WOS against the cyclical unemployment rate. Note that the cyclical unemployment rate corrects for cross-country structural differences in unemployment by subtracting the NAIRU (non-accelerating inflation rate of unemployment) from the prevailing unemployment rate. It can be seen that the self-reported degree of effort at work ("How Hard Work") is negatively correlated with the cyclical unemployment rate across countries. The correlation between the self-reported degree of exhaustion after work and the cyclical unemployment rate does not emerge clearly from the cross-country graph and might be influenced by

Figure 1: Self-reported work effort and cyclical unemployment, cross-country

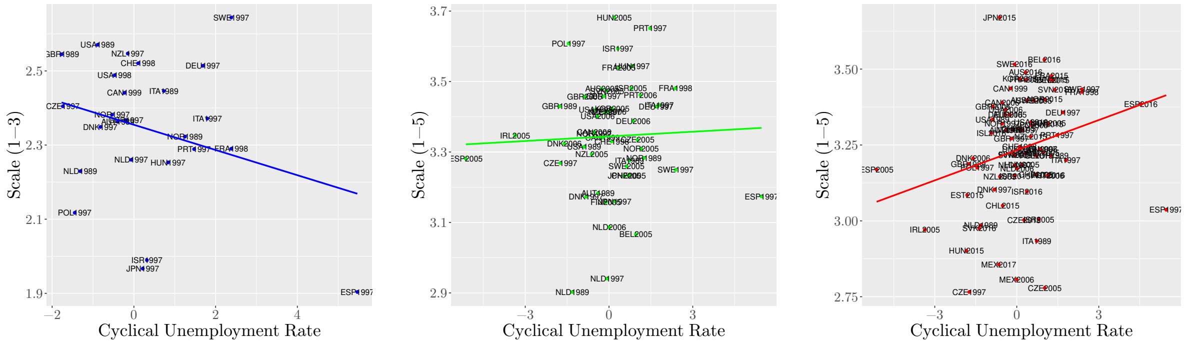

(a) How Hard Work (b) Exhaustion (c) Stress

a few outliers. Self-reported work stress, however, correlates positively with the cyclical unemployment rate. As a first impression, thus, it seems that, when the cyclical unemployment rate rises, people tend to exert less effort on their job, even though they experience it as more stressful. If people are more afraid of becoming unemployed when unemployment is higher, this would support the shirking model. Nonetheless, the fact that people say they work less hard when the cyclical unemployment rate increases suggests that, if the shirking model plays a role in determining the cyclicality of effort, its effect is dominated by the effect of labor hoarding. To trace this suggestive evidence out in more detail, we now turn to our individual-level evidence.

4.2 Individual-Level Evidence

The country-level evidence presented above is suggestive of the labor hoarding view dominating at the macroeconomic level. It cannot, however, be conclusive, because the bivariate correlations presented might obscure more subtle underlying relationships. To take into account additional potential determinants of labor effort, we need to conduct individual-level regressions as well. This section presents the results of this exercise. Two control variables are particularly important: hours and income.

- *Hours.* Controlling for individual hours worked is crucial, since we would like to analyse the cyclicality of effort per hour worked. This is the key labor input we are concerned with, rather than working hours *per se*, which are observable. In the regressions on self-reported effort (WOS), we can control for "number of hours (usually) worked weekly" at an individual level. Note that both the formulation of the question and the inclusion of an answer option, "can't say, varies too much", indicate that this variable concerns actual and not contractual hours. In the regressions on attitudes to effort hours worked are not measured at the micro level, so, there, the control captures average annual hours worked per person employed.

- *Income.* The shirking hypothesis is one particular incarnation of the efficiency wage theory, which says that setting a wage above the reservation wage is optimal for a firm that wants to raise workers' productivity. Effort is a function not only of the unemployment rate, but also of the wage. Therefore, when testing the shirking model, we need to control for wage income, too.

Moreover, we want to control for individual characteristics of the respondent that are likely to influence individual effort at work:

- *Union.* To the extent that firing costs are greater for unionized workers, we would expect labor hoarding practices to be more prevalent.
- *Education.* Workers with higher education might exhibit more pro-cyclical effort if the cost to firms of training and replacing such workers is greater (Blatter, Muehlemann, and Schenker, 2012).
- *Female.* Given that male unemployment tends to rise more than women's during recessions (Hoynes, Miller, and Schaller, 2012), this could bias the unemployment coefficient.
- *Married.* On the one hand, the disciplining effect of unemployment may be greater if the respondent's spouse is dependent on the breadwinner's income (see also next bullet point). One the other hand, the disciplining effect might be weaker for married individuals if they are more financially stable due to their spouse also earning an income. Moreover, if more people divorce in recessions (Hellerstein, Morrill, and Zou, 2013), and married individuals exert either more or less effort on average, this could again bias the coefficient on the unemployment rate.
- *Kids.* The responsibility of supporting a family may also affect the cyclicality of individual work effort if this responsibility makes one more scared to lose one's job in a recession and thus more eager to reduce one's shirking behavior.
- *Rural.* The cyclical nature of effort might be different for workers living in rural and urban areas. It is not *a priori* clear which way the difference will go: rural jobs might be more often industry jobs, which could cause effort to be more countercyclical; yet those who live in rural areas do not necessarily work there, but might simply commute to urban areas.
- *Age.* Younger workers might face a greater risk of losing their job during recessions (Hoynes et al., 2012) if employment protection rises with tenure ("last in, first out").

The precise description of these different individual-level controls can be found in Table 1. We further discuss individual heterogeneity in the cyclical response of labor effort in Section 4.4 below.

Two additional factors that might determine a worker's effort while also being correlated with the unemployment rate are the worker's health and immigration background. Data on these measures are, however, not consistently available in our samples, so we cannot control for them. While this does affect our ability to attribute the procyclical nature of effort solely to the countervailing forces of shirking and labor hoarding, it should not matter too much for our overall estimate of the direction of the cyclicality. In other words, if the cycle does affect labor effort through individuals' health and through migration, we actually want to capture these effects in the coefficient on unemployment. However, not controlling for these factors also impairs our ability to attribute the estimated cyclical movement of effort purely to labor hoarding and shirking effects. As a matter of fact, this double-edged sword holds for most of our other controls as well. Yet, as we also discuss below, inasmuch as the direction of the coefficient on unemployment remains the same whether we include controls or not, our findings regarding the cyclical nature of effort are robust.

Besides the individual-level controls described above, we control for country- and year-fixed effects in all our regressions to account for cross-country differences as well as common time trends. Our use of the cyclical unemployment rate as a baseline measure for the cycle is motivated by the fact that it allows us to control, in addition to the fixed effects, for time-varying structural differences between countries' labor markets. Additionally, the unemployment rate figures either directly or indirectly in the shirking and labor hoarding models and thus constitutes a natural channel through which to estimate their effects.

Self-reported effort at work. Turning to our first set of results, Table 5 reports the regression models for the same WOS effort measures as presented in the cross-country evidence: (1)-(2) how hard the employee works on her job; (3) how often the employee comes home from her job exhausted; (4) how often the employee finds her job stressful, and this only for full-time workers.

— [insert Table 5 here] —

We first report the model with only fixed effects and the unemployment rate for the measure that most closely captures actual work effort, "How Hard Work", in column (1).[6] The coefficient on cyclical unemployment in this stripped-down model is much smaller than 1, with a 1 percentage point increase in the cyclical unemployment rate leading to a sizeable 50% decrease in the odds of people reporting they work hard. Adding individual-level controls to the regression in column (2) retains the sign of this effect and slightly increases its magnitude. The coefficient on cyclical unemployment for exhaustion after work in column (3) is much smaller than the effect on self-reported effort, but also suggests procyclical effort. On the other hand, the same coefficient in column (4) for work-related stress suggests slightly countercyclical stress. Nonetheless, in our results using regional data and country-year fixed effects the coefficients are reversed. The picture that emerges, then, is one of procyclical effort at work. While the coefficients on the control variables do not speak to the cyclicality of effort, as they give average effects, they are interesting in themselves, so we discuss them briefly. Intuitively, an increase in the number of hours usually worked weekly increases all measures of effort reported, by approximately the same amount: a one-hour increase leads to an approximate 5% increase in the probability that the respondent reports a higher level of effort. The lack of significance for the coefficients on hours^2 indicates that self-reported effort is linear in hours worked, although the direction of the coefficients - they are smaller than 1, but due to the large range of hours2, are rounded up to 1 - intuitively points towards a declining slope coefficient on hours, in line with the findings of (Burda et al., 2019). Union membership seems to be associated with higher exhaustion and stress, probably because of a composition effect, where those who experience higher exhaustion and stress are more likely to join a union. When it comes to income, those in the lowest income category report higher self-reported effort for all measures, compared to the median category. On the other hand, those in the above-median categories report to work less hard, even

[6]Recall that the cyclical unemployment rate represents the deviation of monthly unemployment from the NAIRU. Since the WOS sample contained start and end dates of the survey fieldwork in each country, we were able to match this cyclical unemployment rate to the exact period in which the fieldwork took place, which on average had a length of about three months.

though they also report significantly higher exhaustion and stress. This might be due to a difference in relative standards: those in higher-paying jobs might, for the same intensity of work, perceive themselves to be working less hard than those in lower-paying jobs because their standards about what constitutes hard work differ. For education, the average effects are not significant. On the other hand, female workers report significantly higher effort across all measures. This finding can be seen as a sort of 'reverse gender pay gap': whereas the gender pay gap points to lower wages for women who deliver the same work as their male peers, our findings indicate that women who earn the same wage as their male peers - captured by the controls for income - on average report to work harder. Married individuals do not report significantly different effort levels, while those who live in rural areas report significantly lower work-related stress. The effect of the respondents' age on effort exhibits a similar pattern as income: while older respondents report to work increasingly harder than the youngest group (18-25), the effect of age on exhaustion and stress is more hump-shaped. Workers until around 44 years old exhibit increasingly higher exhaustion and stress levels, but after that the coefficient starts falling, with the oldest workers exhibiting lower levels than the 18-25 group. Overall, the average effects draw a demographic picture of work effort that is in line with intuition.

We re-estimate the benchmark model with fixed effects and individual-level controls for various cyclical indicators in Table 6. The HP-filtered unemployment rate represents the deviation of the monthly unemployment rate from its HP-filtered trend, averaged over the period in which the survey fieldwork took place. While we believe that the NAIRU-adjusted unemployment rate is a better measure of cyclical unemployment, we include the HP-filtered rate as a robustness check.

— [insert Table 6 here] —

The estimated coefficient on "How Hard Work" in the regression using the HP-filtered unemployment rate is very similar to the benchmark estimate, but now the effect on exhaustion is larger and more significant. The coefficient on stress retains the same sign and remains significant, although it decreases by about one half.

As an alternative measure of the cycle, we also consider the output gap. Note that we switched the sign of the output gap from its normal definition to facilitate coefficient interpretation. Though the output gap is an often-used cyclical measure, we can only match it to the individual-level variables as a year average. That is, our output gap variable reflects the output gap that prevailed throughout the whole year the survey fieldwork was undertaken. Nonetheless, the coefficients on the output gap are similar in direction to those on the cyclical unemployment rate. In a recession, a decrease in the output gap as it is normally defined, $Y - Y^*$, leads to a decrease in how hard people say they work. Moreover, a decrease in the output gap also leads to a significant decrease in exhaustion due to work. The coefficient on work-related stress is not significant.

Regional unemployment and interacted fixed effects. To subject our main result of the procyclical nature of self-reported effort to a robustness test, we include additional controls for country-year fixed effects, as well as occupation-year fixed effects, in our baseline regression. To be able to do so, we match each respondent's region of residence in the WOS to the unemployment rate prevailing in that region in the year the fieldwork took place. For that purpose, we make use of the OECD regional database, which contains detailed regional economic data going back to 1990. In total, our sample contains 341 regions. These are for the most part distinct regions, despite some regional re-shuffling taking place between the second and the last wave.[7] Besides adding country- and occupation-year fixed effects, we move the unit-fixed effects down from the country level to the regional level. Finally, we exclude the first wave from our sample because there is no regional data available for 1989. This also implies that we can only estimate the coefficient for exhaustion and stress, because the variable for self-reported effort spans only the first two waves.

The outcome of this exercise is reported in Table 7. In contrast to our earlier results, we now report non-cyclical unemployment as well. The reason for this is; on the one hand, the fact that reliable estimates of regional NAIRUs are not readily available; on the other hand, the consideration that the combination of region, year and country-year fixed effects should control for most structural differences across time and space, thus obviating the need for correcting the unemployment rate for such differences by means of the NAIRU. Thus, the coefficients on the regional unemployment rate $U_{r,t}$ can be understood as capturing the effect of a change in the deviation of regional unemployment $D_{r,t}$ from the country-wide average $\bar{U}_t$,

$$U_{r,t} = \bar{U}_t + D_{r,t}. \tag{4}$$

The implicit assumption behind these estimates thus is that the coefficient on $\bar{U}_t$ is similar to the one on $D_{r,t}$. This is reasonable insofar as individuals only care about regional unemployment and not about the country-wide average unemployment, which would be the case if there is limited factor mobility. For the European Union, the empirical literature shows that labor mobility between regions within a given country indeed plays a minor role; for the US, labor mobility between states is more important (Dijkstra and Gakova, 2008). Nonetheless, given that we matched the regional data at the highest regional subdivisions, it seems that factor mobility will indeed be limited. For example, in the US, the regions included are not states but Census Bureau Divisions (e.g. East North Central, Pacific), which are clusters of several states. Finally, despite the fact that non-cyclical regional unemployment should thus be a good measure of the effect of cyclical changes in unemployment, we also report estimates of the coefficients on HP-filtered regional unemployment in Table 7. One would expect the coefficients on this measure to be smaller in magnitude, but not majorly so, since most of the structural unemployment effects should already be filtered out, as explained above.

— [insert Table 7 here] —

The estimates provide robust support for our finding that labor effort is procyclical, since they are all highly significant and smaller than 1. The coefficients for work-related stress have now flipped, but this should not be too surprising, as they were not extremely significant, nor very large in magnitude before. Moreover, since stress seems to be a more subjective experience than exhaustion or self-assessed effort, it was more likely to be related to cyclical shifts in non-economic factors such as survey measurement error or work attitudes. The bias from these can run in either direction: survey measurement error in explanatory variables can lead to attenuation in coefficient estimates, thus reducing the magnitude of the estimates; while higher unemployment could for example lead to respondents being in a worse mood when filling out the questionnaire, thus increasing the magnitude of the estimates insofar as this leads respondents to answer more negatively

[7]Note that due to the inclusion of region-fixed effects, any region that appears in only one wave is effectively dropped from the sample.

to effort-related questions. Indeed, the direction of the increase in the estimates in Table 7 is not clear-cut. However, all coefficients are strongly significant now, as well as being closely in line: both coefficients on regional unemployment estimate an increase of around 9% in the probability of reporting higher effort for a 1 p.p. increase in unemployment, while both coefficient for the HP-filted regional unemployment estimate an increase of around 1-4% for the same. Our finding that work-related stress is procyclical might seem to be at odds with the findings from the biomedical literature, which consistently reports a deterioration in mental health during economic crises (for an overview, see Mucci, Giorgi, Roncaioli, Perez, and Arcangeli (2016)). However, this need not be so: whereas most of these studies look at general mental health, our stress measure specifically asks about how often respondents find their work stressful. It might very well be that individuals' overall stress levels increase in recessions due to, for example, increased job insecurity even as their work-induced stress actually decreases. This would certainly be in line with the fact that we find work-related exhaustion and effort at work to be procyclical as well. In this sense, our finding of procyclical work-related stress might form a bridge between the results from the biomedical literature and those from the economic literature, which, as mentioned, find that temporary upturns decrease health outcomes (Ruhm, 2000).

In all, the picture that arises from these and the above results is that effort at work is significantly procyclical. This finding holds for three different self-reported effort measures, across various cyclical indicators. It is consistent with the cyclical effects of labor hoarding dominating those of the shirking model. In the next subsection, we provide evidence for the mechanism behind both of these models. We also document the individual and occupational heterogeneity in more detail.

One interesting side result arises from repeating these same regressions for a dependent variable which captures the frequency with which the respondent engages in hard physical labor at work. Doing so, we find no significant coefficient on any of the cyclical indicators. Given the sample size of 30,000 observations, this is a remarkable finding in itself. It suggests that the intensity of physical labor does not vary over the business cycle, even as self-reported effort does. Taken together with the fact that the degree of exhaustion does seem to vary significantly over the business cycle, this finding provides some support for Uhlig's (2004) hypothesis that cyclical effort at work depends mostly on time spent "loafing", that is, engaging in work-place leisure instead of actual work, rather than on the intensity of the work actually engaged in.

Attitudes to effort at work. We now look more closely at attitudes to effort. Table 8 reports baseline regression results from a logit model, where the dependent variable is a dummy for whether the respondent regards the "opportunity to use initiative on the job as important. Table 9 reports a similar exercise as before, where we check the baseline results for different effort measures and different cyclical indicators. The first three dependent variables all come from the same overall survey question which asks the respondent to choose non-exclusively among several features he/she finds important in a job (1) an opportunity to use initiative, (2) not too much pressure[8] and (3) that you can achieve something, and they are equal to 1 if the respondent picked the feature in question and 0 otherwise. The first three columns are thus estimated as logit models. The dependent variables in the last two columns, on the other hand, are categorical measures of the extent of agreement with two statements: (4) "hard work brings success" and (5)

[8]We switched this variable around to facilitate coefficient interpretation

"those who don't work turn lazy".[9]

— [insert Tables 8 and 9 here] —

As can be seen from Table 9, the coefficients on attitudes to effort are larger than 1 for nearly all cyclical indicators, indicating that such attitudes are robustly countercyclical. The magnitude of the coefficients is fairly small. For example, the largest effect of a 1% increase in cyclical unemployment on any of the attitude measures is a 5% increase in the odds of the respondent mentioning the possibility to achieve something as important in a job. The small size of the estimated effects aligns with the intuition that values tend to be more stable over time than behavior. That is, in addition to moving in the opposite direction over the cycle, attitudes to effort seem to fluctuate less than actual effort. In light of the earlier findings that self-reported effort is procyclical, one possible interpretation of the countercyclicality of attitudes to effort is that people consider effort at work in a more positive light when they exert less of it.

4.3 Evidence for Labor Hoarding

Although our results so far suggest that labor hoarding plays a more important role than the shirking model in determining the cyclicality of effort at work, we cannot entirely exclude that there are other mechanisms contributing to the procyclical nature of effort that are not labor hoarding. One likely additional driver of the cyclicality of effort aside from labor hoarding and shirking is a labor force composition effect, where firms fire their less productive ("lazier") workers in recessions, which would lead to an increase in average effort at work even if individual effort did not change. We cannot disentangle how much each of these various drivers ultimately contributes to the cyclicality of labor effort. We can, however, test for the presence of labor hoarding and shirking by searching whether its theoretical implications hold empirically.

One of these implications, in the case of the labor hoarding model, is that firms should make more active use of the intensive labor margin when hiring and firing costs are high. To test this prediction, we make use of the OECD's employment protection legislation (EPL) index, which measures the strictness of employment protection for individual and collective dismissals for regular contracts in a host of OECD countries. This synthetic index can take any continuous value from 0 (least strict) to 5 (most strict); in the data it ranges from 0.26 to 4.58 (see Table 4). To the extent that stricter employment protection legislation imposes higher hiring and firing costs on firms, we would expect that the procyclicality of effort intensifies with an increase in the EPL index. This prediction obtains neatly in Figure 2.

— [insert Figure 2 here] —

The figure shows the predicted probabilities obtained from our benchmark model for each of the three levels of the "Work How Hard" effort measure, plotted against the cyclical unemployment rate. Recall that a cyclical unemployment rate of -1, for example, means that unemployment is 1 percentage point below the natural rate of unemployment as measured by the NAIRU. That is, to the left of zero on the horizontal axis, the country is in recession; to the right of zero, the country is in a boom. In the figure, each subplot contains three predicted probability lines, corresponding to the first

[9] See Table 2 for a more detailed description of these measures.

three quartiles of the employment protection index. The sold line corresponds to the highest level of employment protection, the dashed line represents less strict employment protection and, lastly, the dotted line corresponds to the lowest quartile of the EPL index. Consider the first subplot on the left, which reports the predicted probability that a respondent reports the lowest effort level. In accordance with our earlier finding of procyclical effort, the probability lines are upward-sloping. Thus, the higher the cyclical unemployment rate, i.e. the deeper is the recession, the higher is the probability that a respondent will report working at the lowest effort level. Of particular interest here is the finding that the stricter the employment protection legislation, the higher is this predicted probability: the solid line is steeper than the dashed one, which in turn is steeper than the dotted line. In other words, when hiring and firing costs are higher due to more stringent employment protection, effort is more procyclical.

In the rightmost subplot, the predicted probability lines are downward-sloping, again consistent with procyclical effort. In addition, stricter employment protection leads to a lower predicted probability of reporting the highest level in recessions, i.e. when cyclical unemployment is high. Lastly, from the middle subplot in Figure 2, we see that the probability that a respondent will report the medium effort level when unemployment is high is increased when the EPL index is low.

To summarize, the results of this exercise support the prediction of the labor hoarding idea that the procyclical movements of effort will be less pronounced when employment protection is less stringent – and hence hiring and firing costs are low.

4.4 Heterogeneity in cyclical labor effort

Individual Heterogeneity. As discussed in Section 4.2, one might expect the cyclicality of effort to differ for individuals with different characteristics. To scrutinize these differences, we interacted regional unemployment with a host of individual-level dummies in the regression from Table 7. Consequently, we obtained the heterogeneous cyclical responses of exhaustion by evaluating the derivative of exhaustion with respect to regional unemployment at each dummy equal to 1. The results of this exercise are reported in Figure 3, which gives the deviation of the interacted coefficient on regional unemployment from the benchmark coefficient on regional unemployment in the regression on exhaustion, as reported in the first cell of Table 7. Note that the original non-interacted controls are still included in the regression, ruling out the possibility that the deviations are due to composition effects. A positive deviation indicates that effort is less procyclical for people with the given characteristic, while a negative deviation indicates more procyclical effort. The error bars give a 95% confidence interval for the deviation.

— [insert Figure 3 here] —

All interacted coefficients (β_{int}) are significantly different from the benchmark coefficient, except for *unionized* workers, whose cyclical response in work-induced exhaustion does not appear to be different from non-unionized workers. This differs from the finding in Burda, Genadek, and Hamermesh (2017), who find countercyclical effort for unionized workers. However, their finding is not significantly different from zero and in that sense consistent with ours.

Looking further at the individual-specific effects, first, we can see that people living in *rural areas* exhibit more procyclical effort. Part of this effect might be due to the fact that occupations are already controlled for, and thus the countercyclical effort associated with industry jobs (see next section) is already filtered out. It is unclear what does explain the more procyclical effort we find for respondents in rural areas.

To assess the difference in the cyclicality of effort for *married* individuals, we add a dummy that captures whether the respondent has any kids. That way, we can better disentangle the effect of being married from the effect of having a family to take care of.[10] Surprisingly, married individuals have more countercyclical effort than the benchmark. Thus, it seems that if the added financial security that comes from being married does not translate into a lesser fear of job loss. Indeed, it might well be the other way around: insofar as there is a significant marital wage premium, married individuals might actually have more to lose from becoming unemployed and therefore reduce their shirking more in recessions (Antonovics and Town, 2004). Having *kids*, on the other hand, leads to more procyclical effort. This is again surprising, given that one would expect the responsibility of sustaining a family to lead to greater fear or job loss. It might be that parents' exhaustion after work decreases in recessions because their spouse has become unemployed and assumes a greater share of the 'burden' of child-rearing. Alternatively, the combined decrease in both partners' exhaustion after work might be self-reinforcing if it leads to a surplus of energy available in the household for child-rearing, and with that a further decrease in pressure at both partners' work. These speculations seem to be confirmed the fact that including an interaction between *married* and *kids* leaves the deviation from the benchmark for *married* unchanged, but absorbs most of the deviation for kids. This indicates that simply having kids does not affect the cyclicality of one's effort much; it is the combination of being married and having kids that does.

As for *income*, we report the deviations from the benchmark coefficient for the lowest and the highest income group. Both of these exhibit more intensely countercyclical exhaustion, potentially because the stakes of losing one's job are higher for both, and thus the shirking model is more operative.

Women, by contrast, exhibit more strongly procyclical exhaustion, potentially because they are less vulnerable to job loss in recessions than men and thus the shirking model is less operative.

Finally, individuals with higher *education* exhibit markedly more procyclical exhaustion. This finding is consistent with the labor hoarding model, insofar as highly-educated workers are more specialized and costly to replace, and thus firms are more inclined to hoard their labor. It is also consistent with the finding from the next section that more white-collar labor exhibits more procyclical effort, to which we turn now.

Occupational heterogeneity. One might also expect the cyclicality of effort at work to differ between occupations. To assess that difference, we interact the occupation-fixed effects with the regional unemployment rate and report the coefficient on each occupation dummy (β_{occ}) for exhaustion and stress on the job at the mean (μ) $\pm$ one standard deviation (σ) of the regional unemployment rate in Figures 4 and 5, respectively.

For this exercise, we grouped the occupations into the sub-major groups of the International Labor Organization's (ILO's) ISCO88 classification.[11] In the figures, the occupational groups are ranked in accordance with their ranking in the ISCO88 codes, which are given so that one can see which major groups each submajor group belongs to based on the first digit. This ranking roughly coincides with a decrease in skill level from the top of the graph to the bottom. The reference group is "Corporate managers".

[10] We did not include the kids dummy before because it is not consistently available in the dataset.

[11] We thank Ganzeboom and Treiman (2011) for making their code freely available.

The circles in the figure mark the occupation-specific effects on exhaustion and stress, for various level of the unemployment rate. A light blue circle with no fill indicates the coefficient at $\mu - \sigma$ of the unemployment rate; a blue, crossed-through circle marks the same at the mean of the unemployment rate; a black, filled circle at $\mu + \sigma$ of the unemployment rate. The whiskers around the dots are the 95% confidence intervals for the interaction estimate obtained as $\beta_{occ} + \beta_{int}\bar{u}$, where $\bar{u}$ is the unemployment rate kept fixed at one of three levels. Note that for most of the points, the confidence intervals are so small that the whiskers and the points overlap. Procyclical effort is consistent with the colors of the dots becoming lighter as one goes from left to right, since this would imply occupation-specific effort goes up when unemployment goes down.

Practically all estimates are significantly different at different levels of unemployment. We can see this in both figures from the fact that, for most occupations, the three dots and the accompanying whiskers are far enough apart from each other.

— [insert Figures 4 and 5 here] —

From Figure 4, the first thing to note with regard to occupation-specific *exhaustion* is that, overall, more blue-collar occupations - that is, those ranked closer to the bottom of the graph - report to be exhausted more strongly than corporate managers, the benchmark, while more white-collar occupations report to be so less strongly. Conversely, occupation-specific *stress* in Figure 5 is in general much lower for more blue-collar occupations than for corporate managers, while for a few white-collar occupations (Life Science and Health; Teaching), it is higher. Another way to see these patterns is to note that the points in Figure 4 more or less form a downward-sloping line, while those in Figure 5 form an upward-sloping line. These patterns make intuitive sense, insofar as exhaustion is a more physical phenomenon that one would expect to arise from hard manual labor, while stress is a more mental phenomenon that one could expect to arise more often in service-oriented jobs.

Looking more closely at the cyclical response of exhaustion and stress for these occupations, another intuitive pattern emerges. The white-collar occupations generally exhibit more procyclical exhaustion, whereas the blue-collar occupations either exhibit less-pronounced procyclical, or even strongly countercyclical exhaustion. E.g. this is the case for ‘Mining, Construction and Transport’ and ‘Skilled Agriculture and Fishery’, where the dots get darker as we move from left to right (see Figure 4). This seems to be in line with the idea that hard manual labor produces a greater disutility for workers than service-oriented labor (Kaplan and Schulhofer-Wohl, 2018). If this is indeed the case, one would expect more shirking on the job in these jobs, and, consequently, also stronger reductions in that shirking when unemployment increases. Alternatively, blue-collar workers might be more vulnerable to unemployment in recessions and thus more susceptible to its disciplining effect in the shirking model (Hoynes et al., 2012). Another possible explanation, at least in the case of skilled agricultural labor, is that unemployment in this sector does not comove positively with unemployment in the wider economy, or that agricultural output and unemployment are less strongly correlated than its non-agricultural equivalents (Da-Rocha and Restuccia, 2006). As a whole, though, the pattern for exhaustion is consistent with both the shirking model and with labor hoarding, where the former is more active in occupations with stronger disutility from labor, while the latter matters more for high-skilled occupations marked by higher degrees of specialization of employees.

A very similar pattern emerges in the occupation-specific cyclical response of stress on Figure 5, where more white-collar occupations exhibit more pronounced procyclical movements in work-related stress, while more blue-collar occupations exhibit either acylicality in work-related stress, or slight countercyclicality (e.g. 'Drivers and Mobile-Plant Operators').

5 Conclusion

Labor productivity, defined as output per hour worked, is procyclical - the US post-1984 being a noteworthy exception. It is important to understand the reasons for this procyclicality, because they give us some clues on how we should model the macroeconomy. Various economists from Okun (1963) to Fernald and Wang (2016) have argued that the utilization of inputs plays an important role. If the factor in question is labor, this implies labor effort varies over the business cycle, consistent with the labor hoarding view of effort.

Effort is notoriously hard to measure. Proxies based on health indicators, workplace accidents, and absenteeism all point to procyclical effort. Time use surveys paint are more mixed picture, but the overall message is, again, that effort varies procyclically. Direct plant- or firm-level evidence is scarce, but a couple of papers actually come to the opposite finding of countercyclical effort, which supports the shirking model of effort.

In this paper, we take a different route and measure both self-reported effort at work and attitudes to effort at work using data from the ISSP's Work Orientations Survey and the World Values Survey. That way, we find that self-reported effort moves strongly procyclically, while attitudes to effort move slightly countercyclically. This suggests that the labor hoarding model plays a larger role in determining the cyclicality of labor effort than the Shapiro-Stiglitz shirking model. We provide support for this conjecture by looking at the effect of employment protection legislation on the cyclicality of effort. Finally, we document heterogeneity in the intensity of that cyclicality across occupations and individuals.

References

Antonovics, K. and R. Town (2004). Are all the good men married? Uncovering the sources of the marital wage premium. *American Economic Review 94*(2), 317–321.

Basu, S. and J. Fernald (2001, January). Why is productivity procyclical? Why do we care? In *New Developments in Productivity Analysis*, NBER Chapters, pp. 225–302. National Bureau of Economic Research, Inc.

Basu, S. and M. S. Kimball (1997, February). Cyclical Productivity with Unobserved Input Variation. NBER Working Papers 5915, National Bureau of Economic Research, Inc.

Biddle, J. E. (2014, May). Retrospectives: The cyclical behavior of labor productivity and the emergence of the labor hoarding concept. *Journal of Economic Perspectives 28*(2), 197–212.

Bils, M. and J.-O. Cho (1994). Cyclical factor utilization. *Journal of Monetary Economics 33*(2), 319–354.

Blatter, M., S. Muehlemann, and S. Schenker (2012). The costs of hiring skilled workers. *European Economic Review 56*(1), 20–35.

Boone, J. and J. C. van Ours (2002, November). Cyclical fluctuations in workplace accidents. CEPR Discussion Papers 3655, C.E.P.R. Discussion Papers.

Brant, R. (1990). Assessing proportionality in the proportional odds model for ordinal logistic regression. *Biometrics*, 1171–1178.

Buch, C. M. and C. Engel (2013). Effort and redistribution: Better cousins than one might have thought. Preprints of the Max Planck Institute for Research on Collective Goods 2012/10, Bonn.

Burda, M. C., K. R. Genadek, and D. S. Hamermesh (2017, June). Non-Work at Work, Unemployment and Labor Productivity. CEPR Discussion Papers 12087, C.E.P.R. Discussion Papers.

Burda, M. C., K. R. Genadek, and D. S. Hamermesh (2019). Unemployment and effort at work.

Burnside, C., M. Eichenbaum, and S. Rebelo (1993, April). Labor hoarding and the business cycle. *Journal of Political Economy 101*(2), 245–273.

Cahuc, P. and S. Carcillo (2011). Is short-time work a good method to keep unemployment down? *Nordic Economic Policy Review 1*(1), 133–165.

Collard, F. and D. de la Croix (2000, January). Gift Exchange and the Business Cycle: The Fair Wage Strikes Back. *Review of Economic Dynamics 3*(1), 166–193.

Da-Rocha, J. M. and D. Restuccia (2006). The role of agriculture in aggregate business cycles. *Review of Economic Dynamics 9*(3), 455–482.

Danthine, J.-P. and A. Kurmann (2004, January). Fair Wages in a New Keynesian Model of the Business Cycle. *Review of Economic Dynamics 7*(1), 107–142.

Dijkstra, L. and Z. Gakova (2008). *Labour mobility between the regions of the EU-27 and a comparison with the USA*. EU.

Doepke, M. and F. Zilibotti (2017). Parenting with style: Altruism and paternalism in intergenerational preference transmission. *Econometrica 85*(5), 1331–1371.

Ehrmann, M. and P. Tzamourani (2012). Memories of high inflation. *European Journal of Political Economy 28*(2), 174–191.

EVS (2015). European Values Study Longitudinal Data File 1981-2008 (EVS 1981-2008). GESIS Data Archive, Cologne. ZA4804 Data file Version 3.0.0.

Fairris, D. (1998). Institutional Change in Shopfloor Governance and the Trajectory of Postwar Injury Rates in U.S. Manufacturing, 1946-1970. *Industrial and Labor Relations Review 51*(2), 187–203.

Fay, J. A. and J. L. Medoff (1985). Labor and output over the business cycle: Some direct evidence. *The American Economic Review 75*(4), 638–655.

Fernald, J. G. and J. C. Wang (2016, October). Why has the cyclicality of productivity changed? What does it mean? *Annual Review of Economics 8*(1), 465–496.

Galí, J. (1999). Technology, employment, and the business cycle; do technology shocks explain aggregate fluctuations? *The American Economic Review 89*(1), 249–271.

Galí, J. and T. van Rens (2014, March). The vanishing procyclicality of labor productivity. CEPR Discussion Papers 9853, C.E.P.R. Discussion Papers.

Ganzeboom, H. and D. Treiman (2011). International stratification and mobility file: Conversion tools. Amsterdam, Netherlands: Department of Social Research Methodology.

Giavazzi, F., F. Schiantarelli, and M. Serafinelli (2013). Attitudes, policies, and work. *Journal of the European Economic Association 11*(6), 1256–1289.

Giupponi, G. and C. Landais (2018). Subsidizing labor hoarding in recessions: The employment & welfare effects of short-time work.

Hellerstein, J. K., M. S. Morrill, and B. Zou (2013). Business cycles and divorce: Evidence from microdata. *Economics Letters 118*(1), 68–70.

Hoynes, H., D. L. Miller, and J. Schaller (2012). Who suffers during recessions? *Journal of Economic Perspectives 26*(3), 27–48.

Inglehart, R., C. Haerpfer, A. Moreno, C. Welzel, K. Kizilova, J. Diez-Medrano, M. Lagos, P. Norris, E. Ponarin, B. Puranen, et al. (2014). World values survey: Round six-country-pooled datafile version. madrid: Jd systems institute.

Kaplan, G. and S. Schulhofer-Wohl (2018). The changing (dis-)utility of work. *Journal of Economic Perspectives 32*(3), 239–58.

Lazear, E. P., K. L. Shaw, and C. Stanton (2016). Making do with less: Working harder during recessions. *Journal of Labor Economics 34*(S1), 333–360.

Leigh, J. (1985). The effects of unemployment and the business cycle on absenteeism. *Journal of Economics and Business 37*(2), 159–170.

McCullagh, P. (1980). Regression models for ordinal data. *Journal of the Royal Statistical Society: Series B (Methodological) 42*(2), 109–127.

Mucci, N., G. Giorgi, M. Roncaioli, J. F. Perez, and G. Arcangeli (2016). The correlation between stress and economic crisis: a systematic review. *Neuropsychiatric disease and treatment 12*, 983.

Ohanian, L. E. and A. Raffo (2012). Aggregate hours worked in OECD countries: New measurement and implications for business cycles. *Journal of Monetary Economics 59*(1), 40–56. Carnegie-NYU-Rochester Conference Series on Public Policy at New York University on April 15-16, 2011.

Oi, W. Y. (1962). Labor as a quasi-fixed factor. *Journal of Political Economy 70*, 538–538.

Okun, A. M. (1963). *Potential GNP: Its Measurement and Significance*. Cowles foundation paper. Cowles Foundation for Research in Economics at Yale University.

Ruhm, C. J. (2000). Are recessions good for your health? *The Quarterly Journal of Economics 115*(2), 617–650.

Ruhm, C. J. (2015). Recessions, healthy no more? *Journal of Health Economics 42*, 17–28.

Schön, M. (2015). Unemployment, sick leave and health. Annual Conference 2015 (Muenster): Economic Development - Theory and Policy 113013, Verein für Socialpolitik / German Economic Association.

Senney, G. T. and L. F. Dunn (2019). The role of work schedules and the macroeconomy on labor effort. *Labour Economics 57*, 23–34.

Shapiro, C. and J. E. Stiglitz (1984). Equilibrium unemployment as a worker discipline device. *The American Economic Review 74*(3), 433–444.

Solow, R. (1964). *Draft of Presidential Address on the Short-run Relation of Employment and Output*. Verlag nicht ermittelbar.

Solow, R. M. (1979). Another possible source of wage stickiness. *Journal of Macroeconomics 1*(1), 79–82.

Taylor, D. E. (1979). Absent workers and lost work hours, May 1978. *Monthly Labor Review 102*(8), 49–53.

Taylor, J. and R. Taylor (2011). Working hard for more money or working hard to make a difference? efficiency wages, public service motivation, and effort. *Review of Public Personnel Administration 31*(1), 67–86.

Uhlig, H. (2004). Do technology shocks lead to a fall in total hours worked? *Journal of the European Economic Association 2*(2-3), 361–371.

Uhlig, H. and Y. Xu (1996). Effort and the cycle. Technical report, CentER Discussion Paper.

WVS (2015). World Values Survey 1981-2014 official aggregate v.20150418, 2015. World Values Survey Association (www.worldvaluessurvey.org).

Table 1: Description of Variables and Data Sources

Variable	Description	Source
Micro		
Age	Dummy for various age groups, reference group: those younger than 25.	WVS, WOS
Income	Dummy for various income categories, reference group: median.	WVS, WOS
Education: Middle	Dummy for whether respondent's highest form of completed education is secondary.	WVS, WOS
Education: High	Dummy for whether respondent's highest form of education is tertiary (completed or not).	WVS, WOS
Female	Dummy for whether respondent is female.	WVS, WOS
Married	Dummy for whether respondent is married or living together as married.	WVS, WOS
Kids	Dummy for whether respondent has any kids.	WVS
Hours	Categorical variable for number of hours normally worked per week, in main job.	WOS
Macro		
Unemployment	Unemployment rate as percent of total labor force.	OECD
Cyclical Unemployment	Unemployment Rate (as defined above) minus Non-Accelerating Inflation Rate of Unemployment (NAIRU), as calculated by OECD based on Kalman filter estimate of the Phillips Curve.	OECD, own calculations
Output Gap	Potential GDP minus actual GDP, as percentage of potential GDP.	OECD
Recession	Dummy for whether economy is in recession, according to "trough method" interpretation of OECD's business cycle turning points, where recession lasts from first month after cycle's peak to first month after cycle's trough.	OECD, own calculations
Hours	Average annual hours per person employed.	PWT v9.1
Employment Protection	Synthetic indicator of regulation on dismissals and use of temporary contracts, as in force on 1st of January of respective year.	OECD
Unemployment Volatility	Variance of either first differences or deviation from HP-filtered trend of monthly harmonized unemployment rate, divided by variance of ratio of OECD's monthly Composite Leading Indicator (CLI) for GDP to its trend.	OECD, own calculations

Table 2: Description of Effort Measures: Survey Questions

Variable	**Question**
Source: WOS	
	Please tick one box for each item below to show how often it applies to your work.)
Exhaustion	How often do you come home from work exhausted? (1:Never, 5: Always)
Stress	How often do you find your work stressful? (1: Never, 5: Always)
	Which of the following statements best describes your feelings about your job?
How Hard Work	In my job...1. I only work as hard as I have to; 2. I work hard, but not so that it interferes with the rest of my life; 3. I make a point of doing the best work I can, even if it sometimes does interfere with the rest of my life.
Source: WVS	
	...tell me which ones you personally think are important in a job?
Pressure	Not too much pressure.
Initiative	An opportunity to use initiative.
Achieve	A job in which you feel you can achieve something.
	How would you place your views on this scale? ...
Work Success	1: Hard work doesn't generally bring success. It's more a matter of luck and connections - 10: In the long run, hard work usually brings a better life.
	Please specify for each of the following statements how strongly you agree or disagree with it! (1: strongly disagree, 5: strongly agree)
No Work Lazy	People who don't work turn lazy.

Table 3: Summary Statistics World Values Survey

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Macro							
Cyclical Unemployment Rate	68,161	−0.082	1.721	−5.883	−1.050	0.602	6.356
HP-Filtered Unemployment Rate	84,460	−0.343	1.018	−3.725	−0.746	0.279	2.509
Output Gap	69,062	−0.602	2.676	−8.666	−2.538	1.413	6.771
Employment Protection (Index)	69,841	2.175	0.805	0.257	1.702	2.679	4.833
Avg Annual Hrs / Prsn Employed	91,630	1,790	227.125	1,363	1,660	1,901	2,637
Micro							
Import Job Initiative	76,005	0.520	0.500	0.000	0.000	1.000	1.000
Import Job No Pressure	75,988	0.639	0.480	0.000	0.000	1.000	1.000
Import Job Achieve	76,079	0.606	0.489	0.000	0.000	1.000	1.000
Work Brings Success	48,830	6.469	2.639	1.000	5.000	9.000	10.000
No Work Lazy	40,619	3.657	1.135	1.000	3.000	5.000	5.000
Income Scale	69,824	5.707	2.368	1.000	4.000	7.000	10.000
Female	95,846	0.409	0.492	0.000	0.000	1.000	1.000
Married	96,264	0.653	0.476	0.000	0.000	1.000	1.000
Education: Middle	92,133	0.382	0.486	0.000	0.000	1.000	1.000
Education: High	92,396	0.304	0.460	0.000	0.000	1.000	1.000
Age	95,469	39.080	12.026	13.000	29.000	48.000	103.000
Health	82,718	2.011	0.823	1.000	1.000	3.000	5.000
Immigration Background	57,279	0.079	0.270	0.000	0.000	0.000	1.000
Survey Weights	96,693	1.008	0.369	0.000	0.890	1.078	10.284

Table 4: Summary Statistics Work Orientations Survey (ISSP)

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Macro							
Cyclical Unemployment Rate	44,316	0.150	1.356	−5.126	−0.531	1.009	5.468
HP-Filtered Unemployment Rate	43,434	−0.176	1.340	−5.573	−0.700	0.708	4.273
Output Gap	45,363	−0.092	1.900	−4.819	−1.053	0.930	5.582
Employment Protection (Index)	30,794	2.166	0.945	0.257	1.595	2.679	4.583
Micro							
How Hard Work	18,133	2.348	0.697	1.000	2.000	3.000	3.000
Exhaustion (after work)	33,383	3.353	0.874	1.000	3.000	4.000	5.000
Stress (work)	46,344	3.231	0.997	1.000	3.000	4.000	5.000
Hours (average weekly)	45,680	43.443	10.090	0.000	39.000	46.000	120.000
Income Scale	46,767	3.233	1.060	1.000	2.000	4.000	5.000
Female	48,085	0.426	0.495	0.000	0.000	1.000	1.000
Married	47,746	0.627	0.484	0.000	0.000	1.000	1.000
Education: Middle	47,219	0.524	0.499	0.000	0.000	1.000	1.000
Education: High	47,219	0.336	0.472	0.000	0.000	1.000	1.000
Age	47,942	40.745	12.622	0.000	31.000	50.000	94.000
Union	45,753	0.351	0.477	0.000	0.000	1.000	1.000
Survey Weights	48,113	0.999	0.358	0	1.0	1	11

Table 5: Determinants of Self-Reported Effort At Work, Benchmark Logit

	Dependent variable:			
	How Hard Work	How Hard Work	Exhaustion	Stress
	(1)	(2)	(3)	(4)
Macro				
Cyclical Unemployment Rate	0.413***	0.361***	0.957***	1.028***
	(0.020)	(0.001)	(0.008)	(0.010)
Work				
Hours		1.050***	1.040***	1.044***
		(0.017)	(0.008)	(0.007)
Hours^2		1.000	1.000	1.000*
		(0.0002)	(0.0001)	(0.0001)
Union		0.990	1.115***	1.239***
		(0.055)	(0.039)	(0.030)
Income				
1		1.121***	1.176***	1.132***
		(0.037)	(0.029)	(0.024)
2		1.008	1.004	1.030
		(0.078)	(0.043)	(0.029)
4		0.839***	1.168***	1.192***
		(0.045)	(0.036)	(0.044)
5		0.880	1.975***	1.561***
		(0.112)	(0.122)	(0.061)
Education				
Middle		0.832	0.962	0.907*
		(0.103)	(0.051)	(0.051)
High		0.914	0.930*	1.002
		(0.066)	(0.041)	(0.051)
Demographics				
Female		1.485***	1.553***	1.178***
		(0.085)	(0.089)	(0.043)
Married		1.036	0.999	1.030
		(0.044)	(0.039)	(0.029)
Rural		1.047	0.971	0.927***
		(0.102)	(0.040)	(0.024)
Age				
25-34		1.287***	1.124***	1.164***
		(0.097)	(0.044)	(0.045)
35-44		1.573***	1.072***	1.212***
		(0.050)	(0.024)	(0.034)
45-54		1.694***	0.996	1.191***
		(0.067)	(0.028)	(0.032)
55-64		1.905***	0.853***	0.998
		(0.161)	(0.038)	(0.028)
> 64		1.322***	0.626***	0.612***
		(0.004)	(0.002)	(0.003)
Country FEs	X	X	X	X
Year FEs	X	X	X	X
Occupation FEs	X	X	X	X
Observations	7388	7388	21418	35139
AIC	13837	13548	51309	93743
Nagelkerke's R^2	0.067	0.114	0.092	0.089
McFadden's R^2	0.031	0.053	0.036	0.032

[1] All models are ordered logit. Reported estimates are odds ratios, calculated as e^{b_j} where b_j is the jth logit coefficient. Reported S.E.s are obtained via the delta method as $e_j^b * se(b_j)$ and are clustered by country. Estimates and S.E.s are survey-weighted.

[2] Reference groups: (1) Income: median, (2) Edu: no or only primary, (3) Age: < 25.

[3] * p<0.1; **p<0.05; ***p<0.01

Table 6: Determinants of Self-Reported Effort At Work, Various Indicators

	Dependent variable:		
	Work How Hard	Exhaustion	Stress
	(1)	(2)	(3)
Cyclical Unemployment Rate	0.361***	0.957***	1.028***
	(0.001)	(0.008)	(0.010)
Observations	7388	21418	35139
McFadden's R^2	0.114	0.092	0.089
Nagelkerke's R^2	0.053	0.036	0.032
Odds LR Test	0	0	0
Brant Test	0.741	0.008	0.065
HP Unemployment Rate	0.390***	0.806***	1.014**
	(0.002)	(0.004)	(0.006)
Observations	7388	20987	34316
McFadden's R^2	0.114	0.093	0.089
Nagelkerke's R^2	0.053	0.036	0.032
Odds LR Test	0	0	0
Brant Test	0.741	0.007	0.087
Output Gap	0.782***	0.958***	0.986*
	(0.003)	(0.011)	(0.008)
Observations	7388	21889	36117
McFadden's R^2	0.114	0.094	0.09
Nagelkerke's R^2	0.053	0.036	0.032
Odds LR Test	0	0	0
Brant Test	0.741	0.009	0.006
Country FEs	X	X	X
Year FEs	X	X	X
Occupation FEs	X	X	X

[1] All columns are ordered logit models. The output gap is reversed from its normal definition as Y^* - Y. Reported estimates are odds ratios, calculated as e^{b_j} where b_j is the jth logit coefficient. Reported S.E.s are obtained via the delta method as $e^b_j * se(b_j)$ and are clustered by country. Estimates and S.E.s are survey-weighted. The Odds LR test gives the p-value for the Likelihood Ratio test of the multinomial logit model against the proportional odds model. The Brant test gives the p-value for the coefficient on the cyclical indicator.

[2] * p<0.1; **p<0.05; ***p<0.01

Table 7: Determinants of Self-Reported Effort At Work, Regional Indicators

	Dependent variable:	
	Exhaustion	Stress
	(1)	(2)
Regional Unemployment	0.912***	0.909***
	(0.003)	(0.002)
Observations	15074	25947
McFadden's R^2	0.119	0.111
Nagelkerke's R^2	0.046	0.04
HP Regional Unemployment	0.962***	0.987***
	(0.004)	(0.004)
Observations	15190	25712
McFadden's R^2	0.119	0.112
Nagelkerke's R^2	0.046	0.041
Region FEs	X	X
Year FEs	X	X
Occupation FEs	X	X
Cty-Year FEs	X	X
Occupation-Year FEs	X	X

[1] All columns are ordered logit models. Reported estimates are odds ratios, calculated as e^{b_j} where b_j is the jth logit coefficient. Reported S.E.s are obtained via the delta method as $e_j^b * se(b_j)$ and are clustered by country. Estimates and S.E.s are survey-weighted. The Odds LR test gives the p-value for the likelihood ratio test of the ordered logit model with only the unemployment rate against the multinomial logit model with the same.

[2] * p<0.1; **p<0.05; ***p<0.01

Table 8: Determinants of Attitudes to Effort At Work, Benchmark Logit

	Dependent variable:	
	Import Job Initiative	
	(1)	(2)
Macro		
Cyclical Unemployment Rate	1.037*** (0.013)	1.023** (0.012)
Hours		1.000 (0.0004)
Income		
1		0.990 (0.041)
2		0.970 (0.030)
3		0.950 (0.041)
4		0.974 (0.033)
5		0.924*** (0.010)
7		1.092*** (0.037)
8		1.114*** (0.032)
9		1.201*** (0.077)
10		1.414*** (0.045)
Education		
Middle		1.170*** (0.037)
High		1.478*** (0.052)
Demographics		
Female		0.866*** (0.025)
Married		0.967** (0.014)
Kids		0.905*** (0.015)
Age		
25-34		1.191*** (0.027)
35-44		1.158*** (0.030)
45-54		1.028 (0.030)
55-64		1.029 (0.028)
> 64		0.979 (0.061)
Constant	1.942*** (0.250)	1.775 (1.334)
Country FEs	X	X
Year FEs	X	X
Profession FEs	X	X
Observations	26093	26093
AIC	34455	34209
McFadden's R^2	0.096	0.109
Nagelkerke's R^2	0.054	0.062

[1] All models are logit, for full-time workers only. Reported estimates are odds ratios, calculated as e^{b_j} where b_j is the j^{th} logit coefficient. Reported S.E.s are obtained via the delta method as $e^{b_j} \cdot se(b_j)$, are bootstrapped (wild) 10 times by drawing from the country clusters. Estimates and standard errors are survey-weighted. The Odds LR test gives the p-value for the likelihood ratio test of the ordered logit model with only the unemployment rate against the multinomial logit model with the same.

[2] Reference groups: Income: median, Education: no or only primary, Age: < 25.

[3] *p<0.1; **p<0.05; ***p<0.01.

Table 9: Determinants of Attitudes to Effort At Work, Various Indicators

	Dependent variable:				
	Initiative	Pressure	Achieve	Work Success	No Work Lazy
	(1)	(2)	(3)	(4)	(5)
Cyclical Unemployment Rate	1.023*	1.017	1.047***	1.038***	1.054***
	(0.013)	(0.021)	(0.013)	(0.014)	(0.008)
Observations	26093	26079	26094	21506	17014
McFadden's R^2	0.109	0.154	0.144	0.105	0.156
Nagelkerke's R^2	0.062	0.092	0.084	0.025	0.055
Odds LR Test				0	0
HP Unemployment Rate	1.162***	0.941	1.083**	1.018	1.180***
	(0.039)	(0.055)	(0.037)	(0.022)	(0.013)
Observations	26093	26079	26094	21506	17014
McFadden's R^2	0.149	0.179	0.18	0.091	0.156
Nagelkerke's R^2	0.085	0.107	0.107	0.021	0.055
Odds LR Test				0	0.005
Output Gap	1.023***	1.007	1.037***	1.012	0.977***
	(0.008)	(0.018)	(0.011)	(0.008)	(0.007)
Observations	35324	35310	35325	23904	18289
McFadden's R^2	0.149	0.179	0.18	0.091	0.156
Nagelkerke's R^2	0.085	0.107	0.107	0.021	0.055
Odds LR Test				0	0.005
Country FEs	X	X	X	X	X
Year FEs	X	X	X	X	X
Profession FEs	X	X	X	X	X

[1] Columns (1)-(3) are logit models. Columns (4)-(5) are ordered logit models. The output gap is reversed from its normal definition as Y^*-Y. Reported estimates are odds ratios, calculated as e^{b_j} where b_j is the j^{th} logit coefficient. Reported standard errors are obtained via the delta method as $e^b_j \cdot se(b_j)$, are bootstrapped 10 times by drawing from the country clusters. Estimates and standard errors are survey-weighted. The Odds LR test gives the p-value for the likelihood ratio test of the ordered logit model with only the unemployment rate against the multinomial logit model with the same.

[2] *p<0.1; **p<0.05; ***p<0.01.

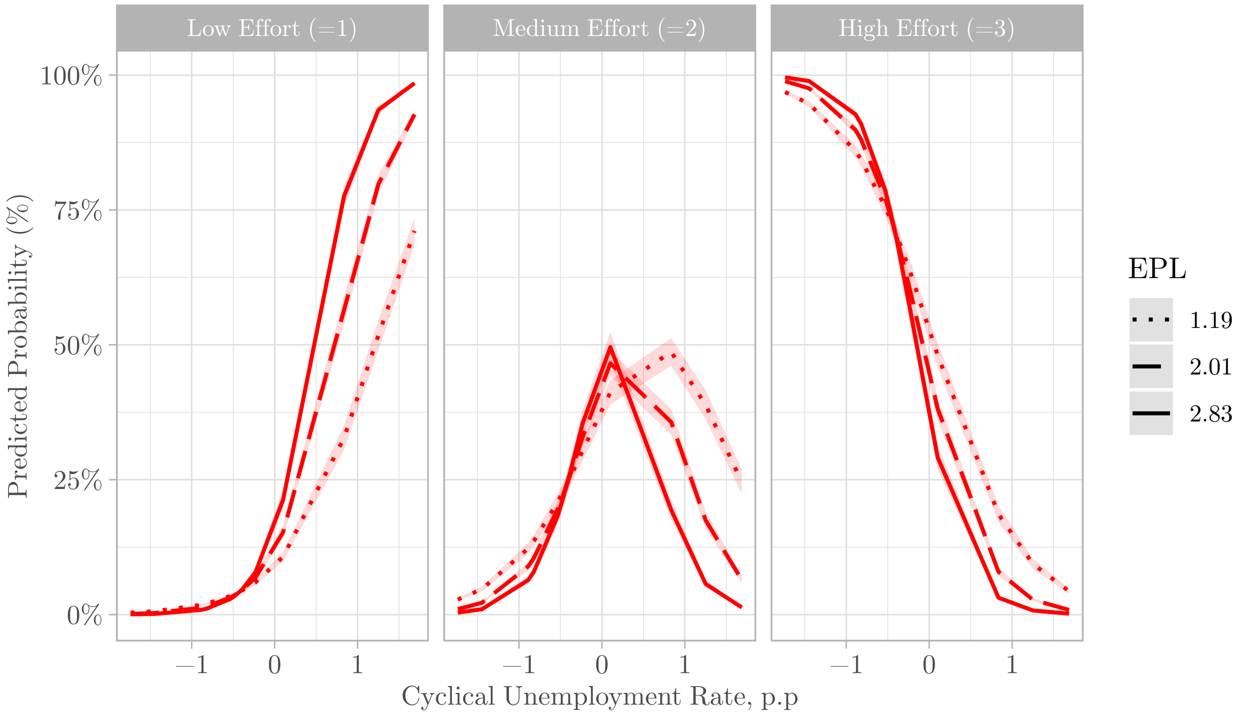

Figure 2: Predicted Probability of Self-Reported Effort Levels vs. Cyclical Unemployment, for Increasing Strictness of Employment Protection Regulation (Q1-Q3), With 95% Confidence Intervals.

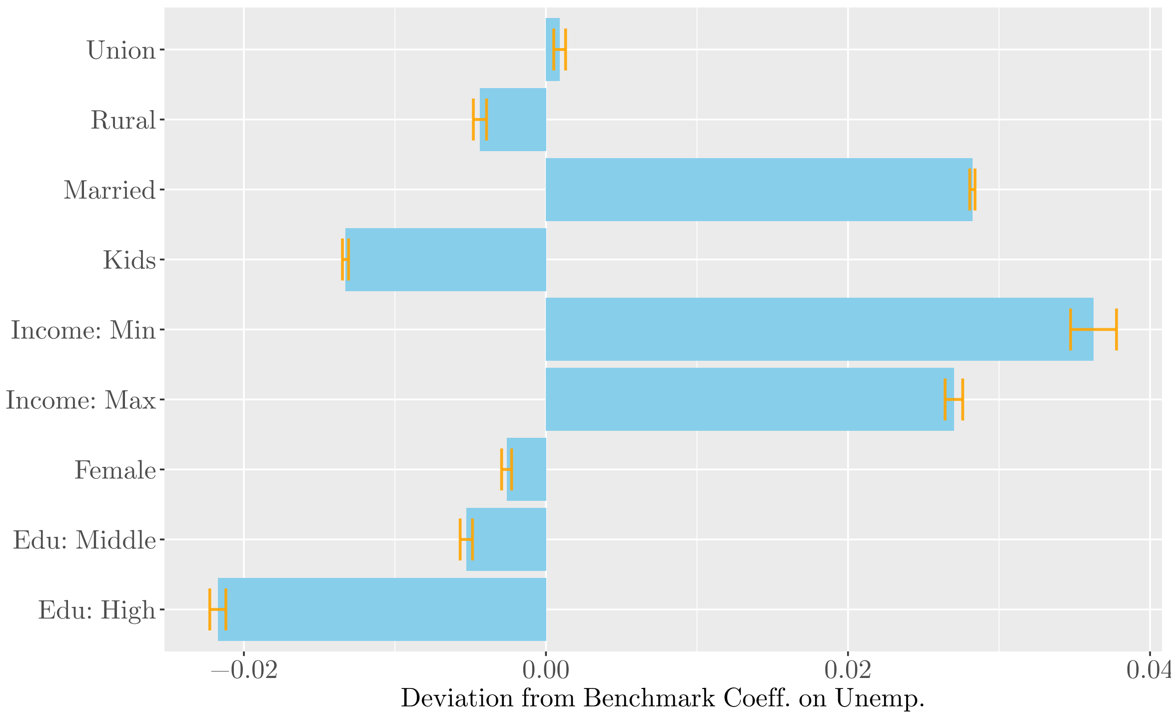

Figure 3: Individual Heterogeneity in Odds Ratio Coefficient for Effect of Regional Unemployment on Exhaustion, with 95% Confidence Intervals.

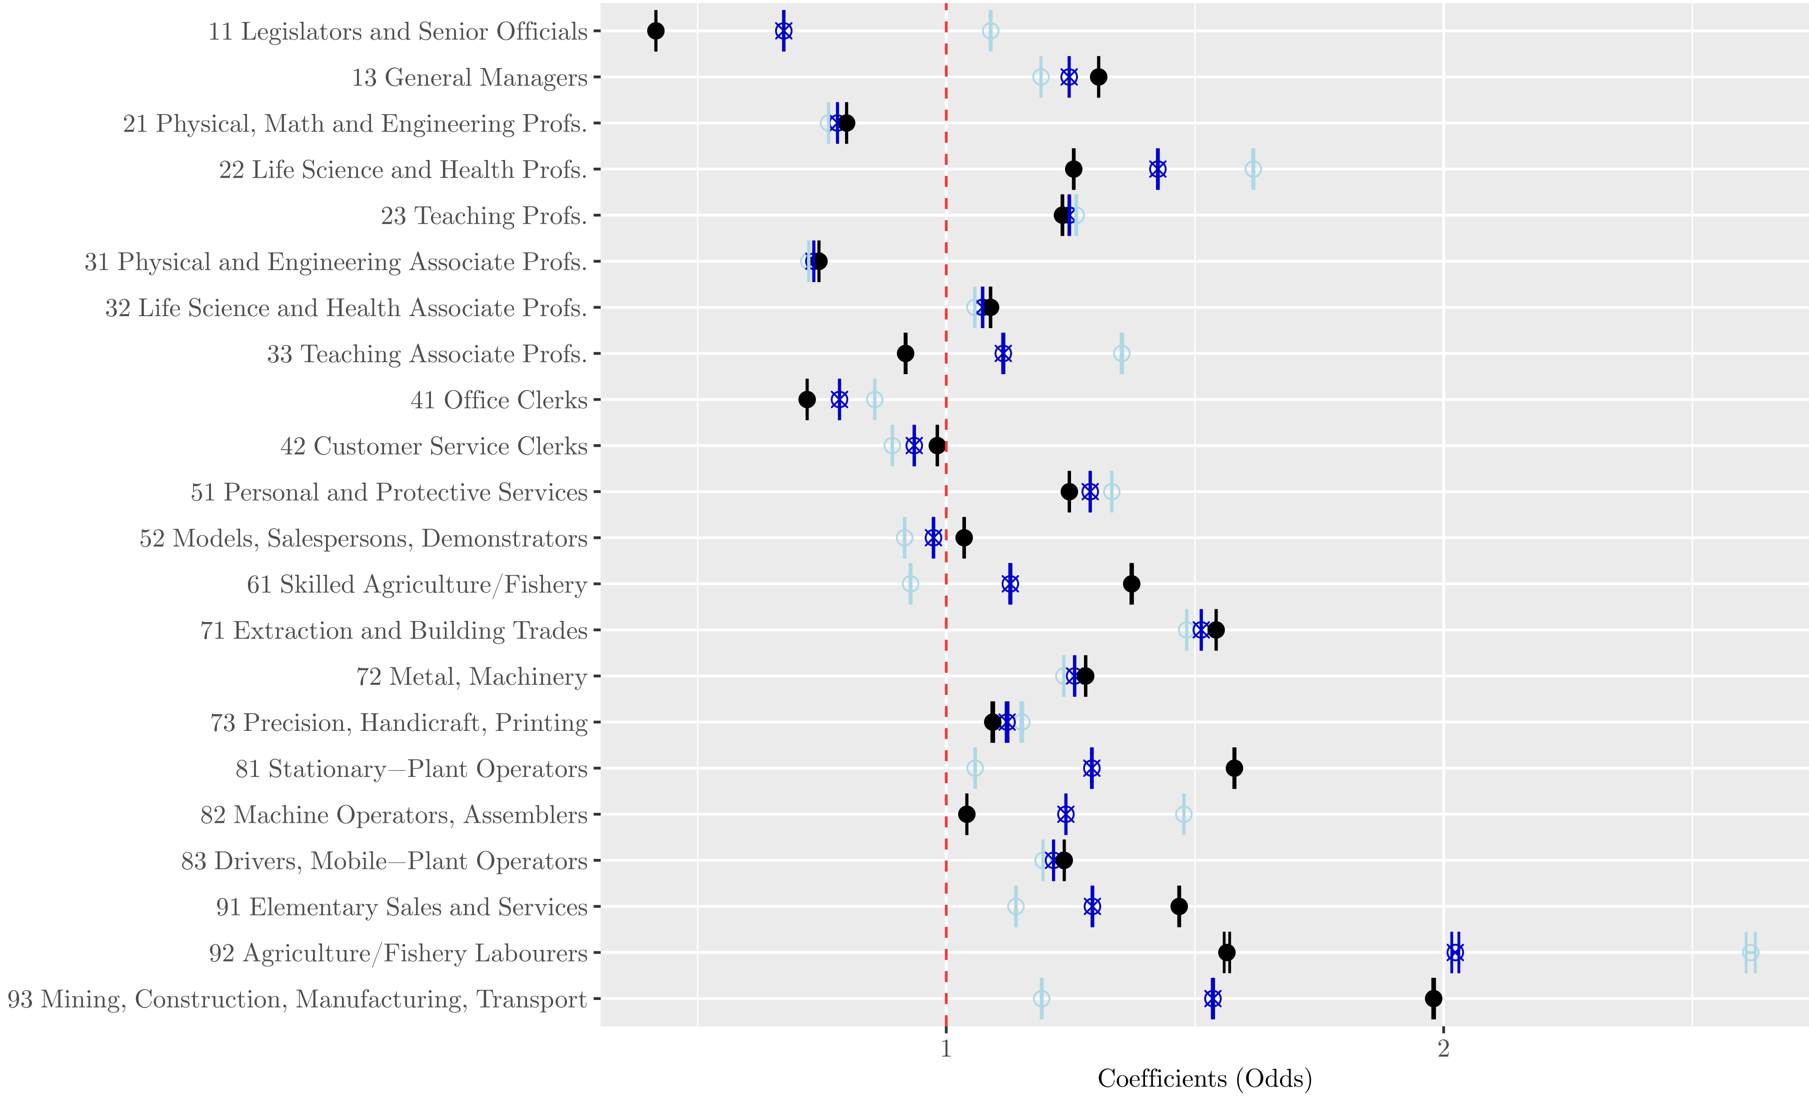

Figure 4: Occupation-Specific Effects on *Exhaustion*, at $\mu \pm \sigma$ of Regional Unemployment, with 95% Confidence Intervals.

Occupational groups are ranked in accordance with the ISCO88 codes. The reference group is "Corporate managers". A light blue circle with no fill indicates the coefficient at $\mu - \sigma$ of the unemployment rate; a blue, crossed-through circle marks the same at the mean of the unemployment rate; a black, filled circle at $\mu + \sigma$ of the unemployment rate. The whiskers give the 95% confidence intervals for the interaction estimate. Note that the confidence intervals are mostly so small that the whiskers and the points overlap. Procyclical effort is consistent with the colors of the circles becoming lighter and losing fill as one goes from left to right.

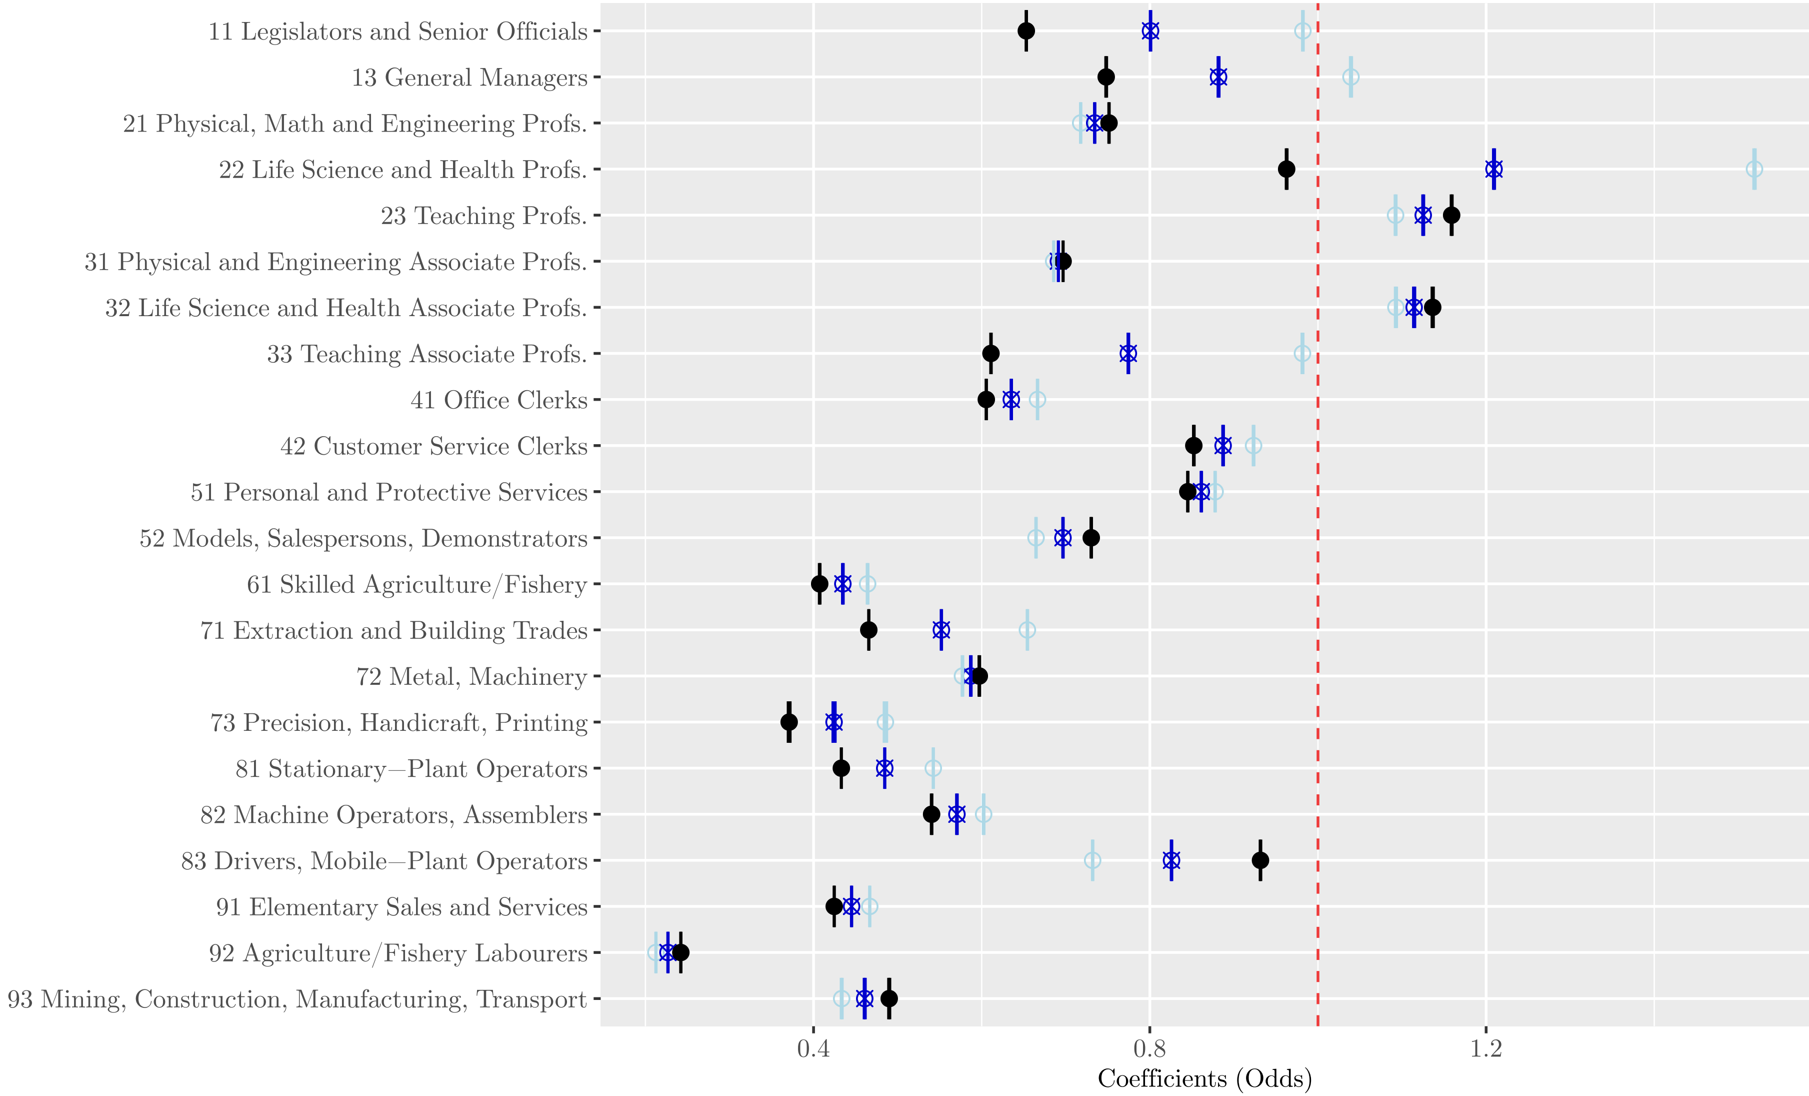

Figure 5: Occupation-Specific Effects on *Stress*, at $\mu \pm \sigma$ of Regional Unemployment, with 95% Confidence Intervals.

Occupational groups are ranked in accordance with the ISCO88 codes. The reference group is "Corporate managers". A light blue circle with no fill indicates the coefficient at $\mu - \sigma$ of the unemployment rate; a blue, crossed-through circle marks the same at the mean of the unemployment rate; a black, filled circle at $\mu + \sigma$ of the unemployment rate. The whiskers give the 95% confidence intervals for the interaction estimate. Note that the confidence intervals are mostly so small that the whiskers and the points overlap. Procyclical effort is consistent with the colors of the circles becoming lighter and losing fill as one goes from left to right.