Remote Control: Aligning Incentives in a World of Remote Work *
WORK IN PROGRESS, PLEASE DO NOT CITE OR CIRCULATE

Pia Andres[1], Carl Benedikt Frey[1,2], and Giorgio Presidente[1,3]

[1]Oxford Martin School, University of Oxford
[2]Oxford Internet Institute, University of Oxford
[3]Bocconi University

June 13, 2023

Abstract

The rise of remote work has created new information asymmetries as fewer workers come into the office, making their actions on the job harder for companies to monitor. In this paper, we examine how firms have responded to this challenge, focusing on the role of high-powered incentives. Building on the Principal-Agent model with moral hazard, which predicts that remote work should increase companies reliance on performance-based pay, we explore this relationship empirically using dictionary methods applied to the near-universe of US online job postings between 2018 and 2022. Controlling for several sources of potential bias and measurement errors, we estimate that remote job vacancies are twice as likely to provide performance pay as onsite ones, making remote workers potentially more vulnerable to adverse economic shocks.

Keywords: remote work, high-powered incentives, performance pay, insurance, monitoring

*We gratefully acknowledge funding from Citi.

1 Introduction

The rise of remote work is an unquestionable feature of the Covid-19 pandemic (Hansen et al. 2023, Aksoy et al. 2022, Adrjan et al. 2021). According to recent estimates, nearly 30% of jobs in the United States are now hybrid or fully remote, and there are good reasons to believe that this shift will persist over the years to come (Barrero et al. 2021). Yet despite a few noteworthy efforts (Bojinov et al. 2021, Choudhury et al. 2021, 2022, Bloom et al. 2015), some key managerial implications of remote work remain poorly understood. In particular, fewer workers coming into the office creates new information asymmetries, as their actions on the job are harder for managers to monitor.

In this paper, we explore how management practises have changed in response, focusing on the role of variable compensation to align incentives. Doing so, we build on Principal-Agent models with moral hazard (Holmström 1979, Mirrlees 1999, Holmstrom & Milgrom 1991), which predicts that remote work should increase the reliance on performance pay as it reduces companies' (the principal's) ability to monitor its employees' (the agent's) efforts, like working hours.

Shedding light on this issue is important, not least since remote work might have unintended consequences, such as compromising the insurance-providing role of the firm (Azariadis 1975, Baily 1974). Indeed, the key idea from the implicit contract literature is that easier access to financial markets allow companies to better diversify risk, making them less risk-averse relative to employees (Rosen 1985). Consequently, companies find it optimal to insulate workers from earnings volatility against lower compensation that discounts the price of the insurance provided. This view has received support in empirical studies, which suggests that firms imperfectly insure workers against shocks (Guiso et al. 2005, Juhn et al. 2018). Undermining the insurance-providing role of the firm does not only potentially make workers more vulnerable to adverse shocks. Given how widespread remote work has become, it could even affect macroeconomic stability more broadly.

For our empirical analysis, we use the near-universe of US online job postings between 2018 and 2022 from Lightcast to capture the incidence of remote work and performance pay using dictionary methods. We tag remote work vacancies based on the list of keywords used by Adrjan et al. (2021) and Hansen et al. (2023), while we develop our own list to tag performance pay vacancies. Our definition of performance pay includes any variable compensation linked to performance, such as bonuses and commissions.[1]

To obtain informative estimates of the correlation between remote work and performance pay, we leverage the granularity of our data and account for the potential presence of omitted variables, such as changing market conditions, economic uncertainty, and other occupational or managerial characteristics. Specifically, we exploit vacancy-level variation to calculate the conditional probability of offering performance pay given that a job is remote, controlling for company-occupation unobserved characteristics.

Motivated by recent research showing that dictionary methods might generate measurement errors when applied to the identification of remote work vacancies (Hansen et al. 2023), we address the possibility of attenuation through an instrumental variable strategy. As is well-known, in response to the pandemic, some U.S. states reacted more vigorously than others in mandating workplace closures. And since the same company often has branches in different U.S. states, we are able to exploit this variation as an instrument for remote work. Indeed, while state-level orders to close workplaces should be correlated to remote work, they are arguably orthogonal to errors in its measurement.

Based on this approach, we find the estimated correlation between remote work and performance pay to be positive and significant across all specifications. Our preferred specification, which accounts for potential time-varying omitted variables and measurement

[1] Thus, our definition of variable pay includes any deviation from fixed base pay. This can be lump-sum and variable bonuses, or bonuses paid at different frequencies.

errors, estimates that remote work vacancies are twice more likely to pay linked to performance, the probability being 0.34. Considering that the share of remote jobs increased by roughly ten percentage points since the pandemics, our estimates imply that remote work increased by 3.4 percentage points or by roughly 25% the share of performance pay jobs—a sizable impact.

Our paper contributes to two strands of the literature. First, we add to an emerging literature examining the rise of remote work and its economic consequences.[2] While survey data shows that workers and employers value flexible working arrangements (e.g. Aksoy et al. 2022, Criscuolo et al. 2021), and results from field experiments suggest that remote work might entail productivity gains (Bloom et al. 2015, 2022, Choudhury et al. 2021, 2022, Bojinov et al. 2021), we know surprisingly little about how remote work has affected workers' compensation. While Barrero et al. (2022) and Brinatti et al. (2021) have examined the relationship between remote work and wages, this paper concerns the structure of compensation and focuses on how remote work affects the use of high-powered incentives by companies.

Second, our paper is related to studies bringing the Principal-Agent model to the data (e.g. Aggarwal & Samwick 1999, Jensen & Murphy 1990, Lazear 2000), as well as the empirical implicit contract literature, pointing to the insurance-providing role of the firm (e.g. Guiso et al. 2005, Juhn et al. 2018). We add to this strand of research by re-examining some of the key issues of these literature's in the context of a contemporary managerial challenge: the rise of remote work. Doing so, we make two contributions. The first one is measurement: to the best of our knowledge, we are the first to capture companies reliance on performance pay across the U.S. economy using job postings data. The second is proposing an application of the multi-task Principal-Agent model of Holmstrom & Milgrom

[2] See also Dalton et al. (2022), Kwan & Matthies (2022), Brynjolfsson et al. (2020), Criscuolo et al. (2021).

(1991) to model remote work, which we view as an exogenous expansion of the set of non-business activities employees can engage in when they work from home.

The remainder of the paper is structured as follows. Section 2 presents our theoretical framework. Section 3 discusses our sources and data construction. Section 4 introduces our econometric approach and presents our results. Section 5 presents a host of robustness tests. Finally, in Section 6, we outline our conclusions.

2 Optimal Contracts When People Work From Home

This section examines how remote work affects the structure of employment contracts through the lens of the Principal-Agent model with moral hazard. In the spirit of Holmström (1979) and Mirrlees (1999), we begin with the most general version of the model, and then sharpen our focus on remote work using the multi-tasking model of Holmstrom & Milgrom (1991). All derivations are provided in Online Appendix A.

Our setup is straightforward. There is a company (the principal) offering an employment contract to an employee (the agent). The employee's utility function is $H(s,e) = U(s) - c(e)$, where s is the compensation she receives, and U is a concave utility function, so that the employee is risk-averse. In our model, $e \in R^+$ denotes effort, and $c(\cdot)$ is an increasing and convex cost function. Moreover, the employee has reservation utility $\underline{\mathrm{H}}$, which represents the minimum amount that she will accept for accepting the employment contract. In our context, we can think of $\underline{\mathrm{H}}$ in terms of labor market conditions, which will be low in bad times and high in good times, when there are many alternative employment opportunities.

Furthermore, the amount of effort spent by the employee affects her performance x, such that $x(e,\theta) : R^+ \times R \to R$, where $\theta \in R$ represents the state of nature outside the control of company or the employee. We assume that more effort leads to better performance, meaning that $\frac{\partial x}{\partial e} > 0$, all else equal. Letting Ω be the set of observable and contractible events, an employment contract is a mapping $s : \Omega \rightarrow R$ specifying the employee's compensation.[3] We also assume that the company is risk neutral and pursues performance minus costs, so that its utility function is given by $\pi(x, s) = x - s$.

This setup generates a dynamic game where the company offers a contract to the agent, and the employee accepts or rejects the contract. If she rejects the contract, the agent receives her reservation utility, but if she accepts it, the agent selects effort e. Alternatively, nature draws θ, determining performance, and the employee receives the payment specified by the contract.

2.1 The General Single-Task Model

We begin our analysis with the most general version of the Principal-Agent model with one task. If we assume that the company can perfectly monitor employees when they are at the office, but not elsewhere, the full-information benchmark breaks down when they work from home. Remote work thus introduces a problem of moral hazard, as the actions (efforts) of the employee is only imperfectly observed when they are not physically at the workplace.

The well-known prediction of this model is that under full information—when, for example, the employee works at the office—the company specifies a desired level of effort—that is, the exact job duties that are expected to maximize profits—and provide a fixed base pay conditional on accomplishing them. However, with incomplete information, such as the inability of observing employees' actions when they work from home, it is optimal to pay at least some variable compensation linked to performance, in order to provide incentives to work hard when they cannot be monitored. This implies that work-from-home job

[3] For simplicity, we assume that the employee has no limited liability.

postings should provide performance pay.

2.2 The Multi-Task Model

Remote work can be modeled more explicitly as an expansion of the set of non-business activities that an employee can engage in while "at work". Specifically, the employee chooses a vector of efforts $e = (e_1, ..., e_n)$ at cost $C(e)$, which yields expected profits for the company $\Pi(e)$. The vector-valued function C is assumed to be strictly convex, while Π is linear. Given partial information, efforts are not directly observable by the company but generate signals $x = t + \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, \Sigma)$ and Σ being the $n \times n$ covariance matrix of the measurement errors.

We assume that the employee has exponential utility with risk-aversion parameter $r \geq 0$, and that the compensation schedule is linear, i.e. $S(x) = w + \alpha x$.[4] The constant term w represents the constant base pay, while αx corresponds to performance-based compensation. We note that in this model, the employee spends a certain amount of effort (or attention) "at work", equal to $E = e + \sum_{a \in A} e_a$. Only e yields a benefit to the company, while A is the set of non-production or leisure activities the employee can engage in, with e_a being the effort allocated to any such activities.

For the employee, the net disutility of effort at work is:

$$D(E) - \sum_{a \in A} v_a(e_a)$$

where $D(\cdot)$ is strictly convex and $v_a(\cdot)$ are strictly concave functions for all $a \in A$. This specification captures two key points. First, attention spent "at work" limits the freedom of the employee and prevents her from spending time on alternative leisure activities. Second,

[4] Holmstrom & Milgrom (1987) show that linear compensation schemes arise optimally in the context of a more general dynamic model.

there are a number of non-production activities that potentially mitigate the disutility cost. Put simply, the option of listening to music or watching YouTube at work can alleviate the fear of missing out on other activities, such as spending time in the park. The set A contains all such non-production activities that are *feasible* for the employee to do. We let N denoting the cardinality of A.

Of course, companies restrict explicitly certain activities by, for example, blocking Facebook on office computers or ban personal phone calls. In addition, while some activities may not be explicitly forbidden, they are nonetheless unfeasible in practice: playing computer games or listening to music during working ours tend to be frown upon, meaning that most people are unlikely to do them at work. But when working from home, such activities are hard to restrict by the company. Building on this intuition, we model the shift to remote work as an exogenous increase in N. This captures the idea that due to companies limited monitoring capabilities at distance, employees experience greater freedom while working from home. Thus, the effort allocated by the employee on the production task is not directly observed by the company, but it generates a signal $x = e + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma^2)$.

In this context, the optimal contract is found by maximising the joint surplus of the company and the employee, subject to the employee's incentive and participation constraints. In order to obtain closed-form solutions, we assume that $D(\cdot)$ and $v_a(\cdot)$ are quadratic, and that $v_a(\cdot) = v_b(\cdot)$ for all $a, b \in A$.[5] Specifically, the optimal performance pay rate is given by:

$$\alpha^* = \left[1 + \frac{r\sigma^2}{1+N}\right]^{-1} \tag{1}$$

[5] The latter assumption is not necessary for our argument, but it allows to make the optimal contract explicitly depending on N.

According to equation (1) an increase in remote work (larger N) should be accompanies by greater reliance on performance-based pay as companies want employees to focus on work-related tasks and prevent substitution to leisure activities. Two considerations are warranted here. First, the extent to which companies rely on performance pay is limited by the uncertainty of the employee's performance, σ^2.[6] The term σ^2 can be interpreted as the variance of a productivity shock affecting the pay-off to the company, or as a measurement error, making it difficult for the company to infer the level of effort exerted by the employee from the realized pay-off.[7] In both cases, uncertainty hampers the incentives to provide performance pay because it makes it either too expensive—the risk-averse employee wants to be insured against low realizations outside her control—or not very useful as observed performance is uninformative about effort.

Second, the multi-task model predicts that remote work increases reliance on performance pay at the intensive margin.[8] Therefore, it allows performance pay even for onsite jobs, with greater reliance on variable compensation for work-from-home jobs. While the distinction is irrelevant for our empirical purposes, as we only observe a binary variable for whether a job is remote or not, it reinforces the case for why we should expect to observe a strong relationship between remote work and performance pay.

3 Data and Measurement

For our empirical analysis, we characterize employment contracts using the near-universe of US online job postings between 2018 and 2022 from Lightcast. Lightcast provides the full text of the job posting, which allows us to tag vacancies offering remote work and

[6] This is a fairly general property of the Principal-Agent model, as discussed in Online Appendix A.2.1.

[7] This equivalence arises because the company is risk-neutral, and so it only cares about the expected value of the pay-off.

[8] This is contrast to the general, single-task model, which only makes a prediction on the extensive margin—either fixed based pay or a combination of fixed and variable component.

performance pay using dictionary methods. Based on this approach, we create a dataset consisting of 146,459,899 vacancies from 3,240,913 employers[9] across 1,061 8-digit occupations.[10]

To create this dataset, we proceed as follows. First, we search for keywords in the description of each vacancy v, and define two dummy variables: $P_v \equiv \mathbb{1}(v \text{ mentions performance pay})$ and $R_v \equiv \mathbb{1}(v \text{ mentions remote work})$. We use the keyword list in Adrjan et al. (2021) and Hansen et al. (2023) to tag remote work vacancies, provided in Table 1. We next minimize the number of false positives by removing vacancies where keywords appear but are mentioned to explicitly exclude the possibility of remote work. Specifically, we "untag" vacancies in which remote work keywords appear in the form of "no * keyword", "keyword * any special character no", "keyword * no", "keyword ***** not an option". These phrases were identified as frequently driving false positives in manual checks.

Turning to performance pay, the keywords used to tag ads are based on manual checks on the texts of job ads and are listed in Table 2.[11] False positives are an issue here as well, as some ads mention sign-on or referral bonuses or qualifications which are considered a "bonus" in a candidate. We deal with these in the same manner as for remote work, namely by "untagging" ads in which the keyword is invalidated by the surrounding text such as in "sign-on bonus", "signing bonus" or "referral bonus". In addition, common sources of false positives are identified via random checks of the texts of ads flagged as including flexible pay. We note that the word "Commission" is sometimes used to describe an institution,

[9]Because firm names in Lightcast are sometimes spelled inconsistently, we harmonize them as far as possible by setting employer names to lower case and removing special characters and legal suffices.

[10] Occupational categories follow the ONET classification, which is similar to the Standard Occupational Classification (SOC) system.

[11] One example of a vacancy tagged as involving both remote work and flexible compensation practices is entitled *Interviewing for Healthcare Benefits* and reads *"Medical Organization is looking for friendly telephone workers to enroll new members into Discount medical programs. Successful Reps will: Earn Top dollar Receive Health & Dental Benefits for the Family Earn Top Dollar!* ***Excellent Commissions & Bonuses*** *No Cold Calling Opportunity for Promotion Requirements: Must Be Able to* ***Work From Your Home Office*** *Telephone, Computer & Internet Access."*

Table 1: Keywords Used to Identify Remote Work Vacancies

100 percent remote	100% remote	location remote
percent remote	remote: yes	remote option
remote workable	remote work teleworking	telework
work * remotely	work from home	work remote
working * remote	working remotely	fully remote
partially remote	remote assignment	remote first
remote position	remote working	telecommute
teleworking	work at home	work from home: yes
work remotely	working from home	home based
location * remote	partly remote	remote based
remote initially	remote work	remote yes
telecommuting	work * remote	work at home: yes
work from home yes	workable remote	working remote
home office		

Notes: This table presents keywords used to identify vacancies as being associated with remote work. * refers to 0 or 1 word of any length.

such as the American Joint Commission on Cancer. To address this, we carry out robustness checks in which ads tagged due to the inclusion of "commission" are dropped.

3.1 Descriptive Evidence

To better understand the incidence of variable compensation, Figure 1 shows the average share of performance pay job postings (over all job postings) by broad occupation category obtained with our procedure. We note that there is considerable variation across occupational categories, with education and library jobs having the lowest share of performance pay with 9.6%, while we find management jobs are in the middle of the distribution at 16%. Unsurprisingly, we find that sales jobs are the most likely to offer performance pay, and do so in almost a third of cases (28%).

Moreover, to see how our measures of remote work and performance pay have evolved

Table 2: Keywords Used to Identify Flexible Remuneration Practices

bonus	stock option
commission	performance based salary
salary ** performance based	performance based pay
pay ** performance based	performance based compensation
compensation ** performance based	performance based wage
wage ** performance based	performance based remuneration
remuneration ** performance based	flexible salary
salary ** flexible	flexible pay
pay ** flexible	flexible compensation
compensation ** flexible	flexible wage
wage ** flexible	variable salary
salary ** variable	variable pay
pay ** variable	variable compensation
compensation ** variable	variable wage
wage ** variable	

Notes: This table presents keywords used to identify vacancies as being associated with performance-based pay. * refers to 0 or 1 word of any length.

over time, Figure 2 plots the shares of performance pay and remote job postings as captured by our dictionary method between 2018 and 2022, where we re-weight the job postings in each quarter to match the 8-digit occupational distribution before 2020. The solid line in Figure 2 represents the share of remote job postings, which increases from 3% pre-pandemic to more than 15% in 2021-Q3, then falling back and stabilizing at around 7%. The dashed line, on the other hand, represents the share of performance pay postings. In 2018-Q1, roughly 15% of job postings reported some form of performance pay. We note that the series follows a slightly increasing trend, gaining 5 percentage points before the last quarter of 2020. Around this time, companies significantly increase their use of performance pay, with the share increasing up to 28% in 2021-Q4, at which point it stabilizes.

We note that while the use of performance pay has increased over the sample period,

Figure 1: Share of performance pay jobs by broad occupation category.

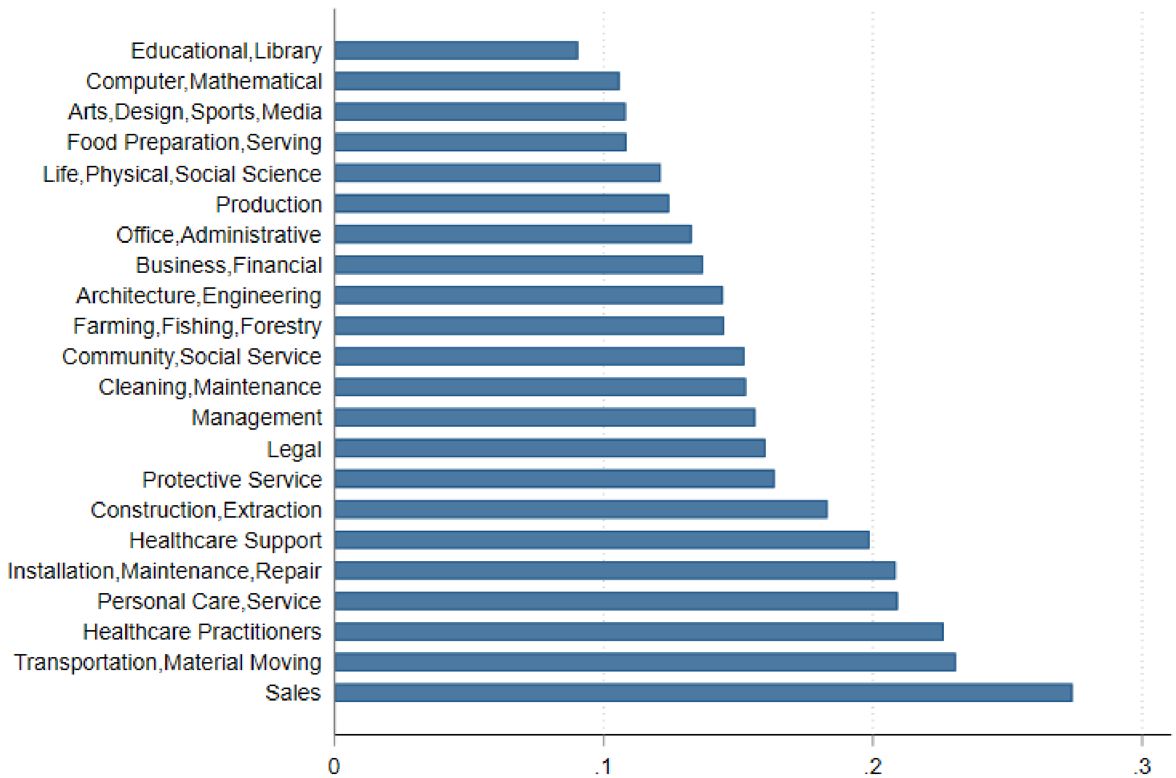

Notes: This figure presents the sample average share of performance pay online job postings by broad occupational category. To tag performance pay vacancies, we use dictionary methods based on the keywords in Table 2.

there does not seem to be an obvious correlation with the share of remote jobs. The largest increase in performance pay, taking place in 2020-Q4, lags the rise of remote work by 3 quarters. One possible explanation is that the pandemic was a source of much uncertainty (Altig et al. 2020). As discussed in Section 2, uncertainty about the economic environment can hamper incentives provision in the Principal-Agent model. In fact, data from the U.S. COVID Uncertainty Index (Baker et al. 2020), plotted in Figure B1 of the Online Appendix, shows that consistently with the theory, the share of performance-linked remote jobs appear to take off only when uncertainty begins to fade. However, the trend observed in performance pay might also be driven by other factors, such as the need to sort workers. For one thing, the literature suggests that companies might offer performance pay to attract employees that are confident about their productivity (Lazear 1986, 2000). We note that also this explanation seems plausible, given the labor market tightness observed by Autor et al. (2023) in the immediate aftermath of the pandemic.

Figure 2: Shares of performance pay and remote job postings.

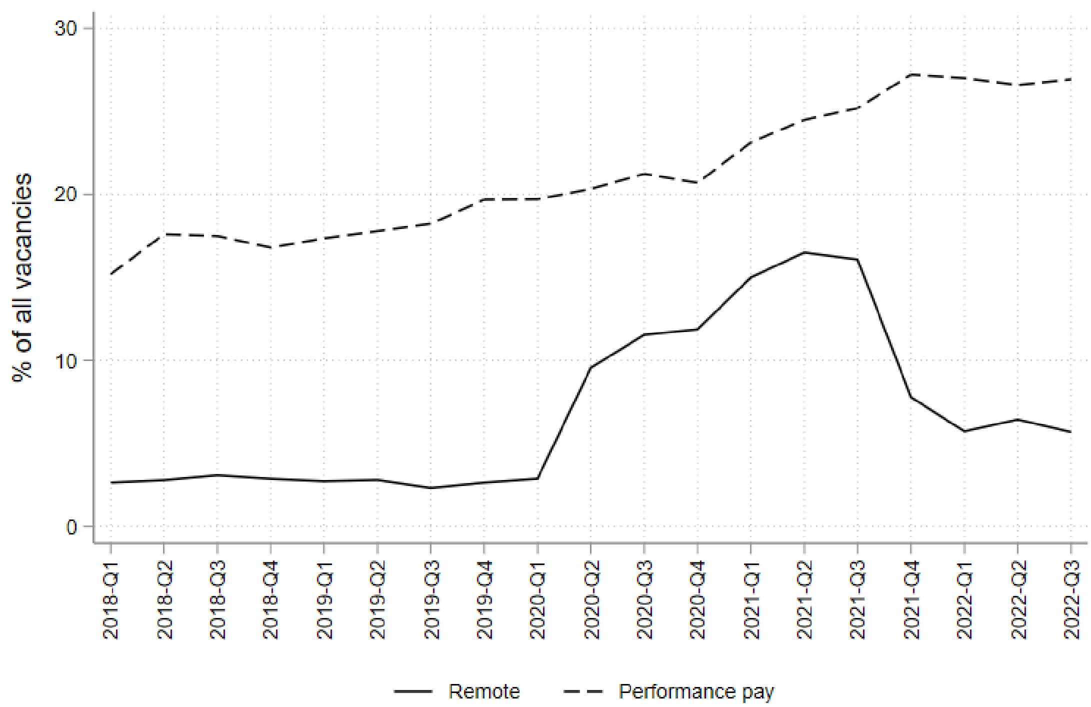

Notes: This figure presents the sample average shares of remote and performance pay job postings by quarter. To tag remote job vacancies, we use dictionary methods based on the keywords in Table 1. To tag performance pay vacancies, we use dictionary methods based on the keywords in Table 2. We re-weight the job postings in each quarter to match the 8-digit occupational distribution before 2020.

Finally, our model, outlined in Section 2, suggests that the relationship between remote work and performance pay might depend on occupation-specific characteristics, such as the difficulty to infer employees' effort conditional on observed performance. To shed light on potential occupation-specific patterns, Figure 3 presents the share of performance pay and remote vacancies by broad occupational category. For all occupation categories performance pay is more likely to be provided in conjunction with remote work, which is consistent with the notion that companies rely more on high-powered incentives when they cannot monitor their employees.

Figure 3: Performance pay provision of onsite *vs* remote job postings by broad occupation category

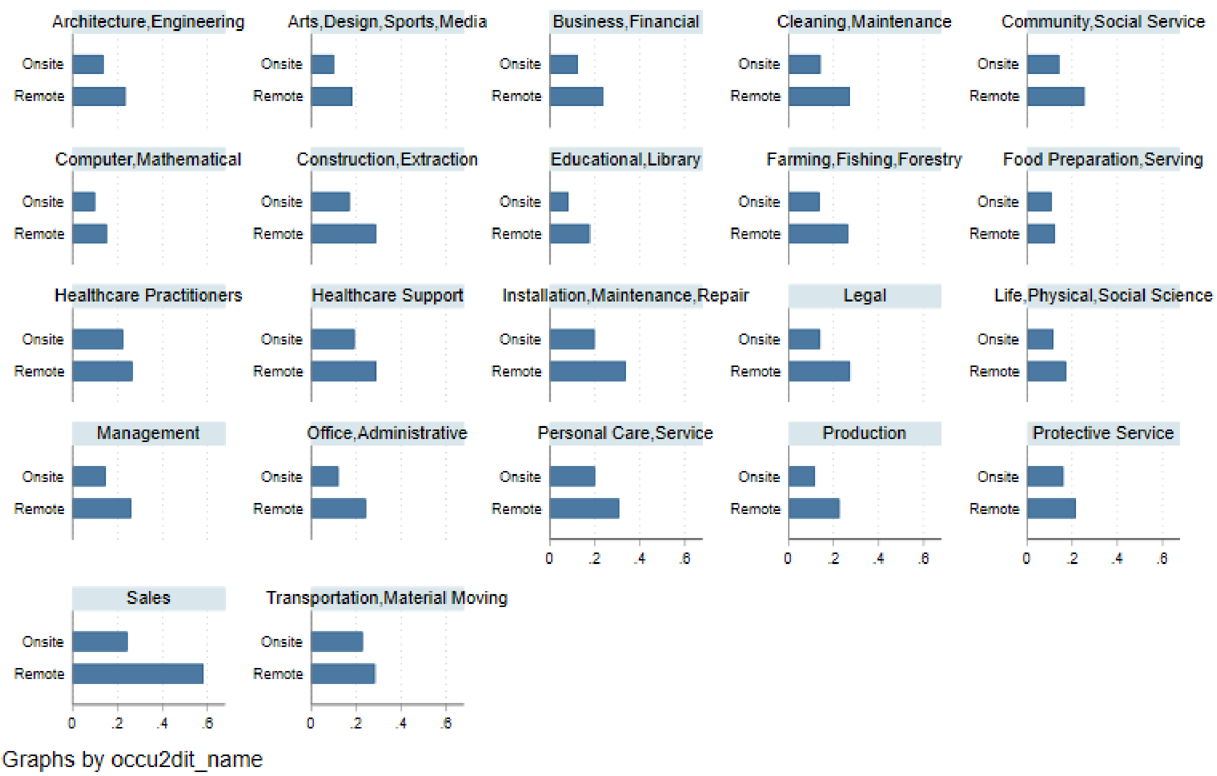

Notes: This table presents the share of performance pay job postings by onsite and remote status by broad occupational category. To tag remote job vacancies, we use dictionary methods based on the keywords in Table 1. To tag performance pay vacancies, we use dictionary methods based on the keywords in Table 2.

4 Measuring the Correlation between Performance Pay and Remote Work

To examine the relationship between remote work and performance pay more systematically, we begin by estimating the the following linear model with OLS:

$$P_v = c + \beta R_v + u_v \tag{2}$$

where c is a constant, P_v and R_v are the dummy variables discussed in Section 3, and u_v is an error structure, which we discuss in the below. Column 1 of Table 3 presents the results of regressing P_v on R_v and a constant term, which represents the average probability of an onsite vacancy providing performance pay.[12] We note that the estimated coefficients are both statistically significant at the 99 percent level and imply that onsite vacancies have a 17% probability of offering performance pay. This probability, however, increases to roughly 30% when the job is remote. We note that on average, over the sample, remote work increases the probability of a job providing compensation linked to performance by 70%, which is line with the predictions of the Principal-Agent model.

However, while the coefficients in column 1 are tightly estimated, the R-squared is lower than 0.01. This speaks to existing evidence, revealing a great deal of heterogeneity in the use of remote work in the United States, both across companies and occupations (Hansen et al. 2023). This variation can arise due to difference in task-feasibility for remote work, attitudes of management, and worker preferences. Since these factors are likely to affect performance pay provision too, the correlation presented in column 1 might simply reflect sample-specific characteristics, meaning that it might not be very informative.

To better understand the dimensions of the data that explain the sample variation in

[12] Unless differently stated, in all specifications we cluster errors at the company-level.

performance pay, we estimate Equation (2) using different sets of fixed effects, which we include in the error term u_v.[13] For example, Column 2 of Table 3 adds month fixed effects, which we note do not substantially alter the estimated correlation relative to column 1. In addition, in column 3, we include company fixed effects, whose inclusion halves the estimated correlation between remote work and performance pay but increases the R-squared to 0.342. In column 4 we include 8-digit occupation fixed effects, which again do not substantially alter the estimated correlation and R-squared relative to column 1. We conclude that a substantial share of the sample variation in performance pay is driven by differences between companies, rather than between occupations or variation over time.

Finally, in column 5, we include company-occupation and month fixed effects, which corresponds to the Two-Way Fixed Effect (TWFE) estimator that is widely used for panel data. The coefficient drops by roughly half relative to column 2, suggesting that the fixed effects absorb additional unobserved variation that is positively related to performance pay provision and remote work. Thus, to minimize omitted variable bias, we use these fixed effects in all the following specifications. The magnitude of the coefficients in column 5 imply that remote vacancies are 5.6 percentage points more likely to offer performance pay, or 30% more likely than onsite vacancies.

4.1 Measurement Errors: 2SLS Estimates

Clearly, measurement errors due to potential fallacies of the dictionary methods used to identify remote work vacancies are a cause of concern (Hansen et al. 2023), not least since this might generate attenuation bias.

To address this issue, we employ an instrumental variable strategy. Our instrument is based on the fact that companies have branches in different U.S. states, and that at

[13] Note that we still include a constant term, because the specifications considered in this sections are pooled cross sections of vacancies with different intercepts depending on the chosen fixed effects.

Table 3: OLS estimates of the vacancy-level relationship between remote work and performance pay.

	(1) Performance pay	(2) Performance pay	(3) Performance pay	(4) Performance pay	(5) Performance pay
Remote job	0.117***	0.107***	0.062***	0.107***	0.056***
	(0.008)	(0.009)	(0.005)	(0.007)	(0.006)
Constant	0.166***	0.167***	0.170***	0.167***	0.171***
	(0.004)	(0.004)	(0.000)	(0.003)	(0.000)
Observations	116,998,714	116,998,714	116,014,879	116,998,714	111,746,792
R-squared	0.007	0.016	0.337	0.070	0.502
FE	none	month	company	occupation	company-occupation and month

Notes: This table presents OLS estimates of Equation (2). To tag remote job vacancies, we use dictionary methods based on the keywords in Table 1. To tag performance pay vacancies, we use dictionary methods based on the keywords in Table 2. Standard errors are clustered at the company-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

the onset of the pandemic, states implemented a staggered mandate of workplace closing rules. Overall, 36 states imposed workplace closing for *all* non-essential services on March 2020, while 4 followed by a lag between one to seven months, and 11 mandated closings of only some sectors or categories of workers.[14] In our sample, 85% of recruiters hire in more than one state, while 7% hire in all states, and recruiters hire in 22 states on average.

After March 2020, remote work increased much more rapidly in mandating states, as shown by Online Appendix Figure B2. It presents the average difference in the share of remote job postings issued by companies in mandating *vs* non-mandating states. Both series increase steeply at the onset of the pandemic and reverse towards zero in the final part of the sample, as companies in non-mandating states increase their reliance on remote work with a lag relative to mandating ones. We exploit this policy-induced variation as an instrument for remote work.

For identification purposes, we drop the 4 laggard states, so that the treatment group

[14] We collect daily data on US state level Covid-19 policy responses from the Oxford COVID-19 Government Response Tracker.

includes companies in states mandating strict workplace closings on March 2020, while the control group includes companies that did not mandate strict closings at all. We cut the sample after 2021-Q3 in order to exploit the exogenous variation due to regulation and avoid contaminating our estimates with the later reversal, and we only keep company-occupation-state units that we observe both before and after March 2020.[15]

Specifically, we implement a 2SLS estimator using the following system of equations:

$$R_{vrst} = u_{rs} + u_t + \tau\, \Psi_s \times \mathbb{1}\{t \geq \text{2020-03}\} + X'_{st}\boldsymbol{\xi} + \epsilon_{vrst} \tag{3}$$

$$P_{vrst} = u_{rs} + u_t + \phi R_{vrst} + X'_{st}\boldsymbol{\xi} + \varepsilon_{vrst} \tag{4}$$

Equation (3) is the first stage regression. The company-occupation (r)-state (s) fixed effect is denoted by u_{rs} and month fixed effect by u_{rt}. The instrument $\Psi_s = 1$ if state s mandates closing in March 2020, and zero otherwise. We define the treatment variable as an absorbing state, which captures the persistent effect of forced policy experimentation and changing attitudes and preferences towards remote work (Barrero et al. 2021). This is consistent with the fact that the share of remote jobs did not revert to the pre-pandemic level as the restrictions were lifted (see Online Appendix Figure B3).

Of course, cross-state differences in policy, which we exploit as our instrument, might be largely attributed to factors such as demographics, geography and politics, which can be considered time-invariant over the relatively short period considered. As such, their impact should be absorbed by the fixed effects. However, since the instrument varies at the state-month level, we also include a vector of state-level controls X_{st}, log-real state-level GDP and the log of COVID-19 cases.[16] The inclusion of these key variables further mitigates

[15] Results are very similar if we use the full sample.

[16] Daily PCR test results by US State are taken from healthdata.gov. Quarterly state-level GDP in millions

the concern of endogeneity of workplace closing policies, which might have reacted to state-specific economic and public health developments.

Equation (4) is the second stage regression. The coefficient of interest is ϕ, which as before measures how much more likely remote vacancies are to offer performance pay—similar to α^* in Equation (1).[17] We note that the fixed effects in (4) should absorb at least part of the confounding variation in r, σ^2 and possibly other confounders due to differences in cost functions across units.

Table 4 presents the results of estimating system (3)-(4).[18] Column 1 shows that the first stage is tightly estimated: companies recruiting in states mandating strict workplace closing rules were more likely to post remote jobs. The coefficient is significant at the 99 percent confidence level and the R-squared is roughly 0.4. In column 2, the second stage coefficient is positive and significant at the 99 percent confidence level and the first stage F statistics is very large, which allows us to reject the hypothesis of weak instrument. The size of the coefficient is almost four times larger than the OLS estimates in column 5 of Table 3, which is consistent with the presence of measurement errors and attenuation bias. The estimated coefficient implies that remote jobs are almost 20 percentage points more likely to provide performance pay than onsite ones.

5 Robustness

To probe the robustness of our results, we begin by experimenting with a more granular set of fixed effects that takes into account the possibility that the same company hires the same occupation but for different areas of business. For example, in our sample, job postings for

of chained 2012 dollars from the US Bureau of Economic Analysis.

[17] However, while α^* in equation (1) captures how much performance pay is optimal to provide, we measure performance pay with a binary variable, and so we can only capture the extensive margin.

[18] We use two-way clustering of errors at the company-occupation and state-date level.

Table 4: 2SLS estimates of the vacancy-level relationship between remote work and performance pay.

	(1) Remote job	(2) Performance pay
Remote job		0.181***
		(0.018)
State mandates workplace closing (yes = 1)	0.011***	
	(0.000)	
Observations	54,397,497	54,397,497
R-squared	0.419	–
Company-occupation FE	yes	yes
Month FE	yes	yes
State-level controls	yes	yes
First stage F	–	5685

Notes: This table presents 2SLS estimates of system (3)-(4). The instrument $\Psi_s = 1$ if state s mandates closing in March 2020, and zero otherwise. The vector X_{st} includes log-real state-level GDP and the log of COVID-19 cases. To tag remote job vacancies, we use dictionary methods based on the keywords in Table 1. To tag performance pay vacancies, we use dictionary methods based on the keywords in Table 2. Standard errors are clustered at the company-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

data scientists from Amazon are sometime classified in 6-digit NAICS "Electronic shopping and mail-order houses", while in others instances they are classified in "Software publishing". While the drawback of this approach is that we lose a large number of observations due to missing industry information, Online Appendix Table B2 shows that the estimated 2SLS coefficient is positive and significant at the 99 confidence level. However, it is more then ten percentage points larger than in the full sample.

In addition, Online Appendix Table B3 shows that our results do not depend on sampling choices. In particular, we document that our results are still valid if, in estimating the system (3)-(4), we do not drop laggard states, exclude recruiters observed only in the pre- and post-pandemic period, and do not cut the sample after 2020-Q3. However, also in this case the estimated correlation is substantially larger than in the baseline.

Finally, so far, we have relied on specifications based on daily vacancy-level information and estimated linear probability models using the dummies R_v and P_v. We note that this approach identifies the correlation based on a repeated cross section of vacancies for each company-occupation pair in each month. While intuitive, this is not the only possible way to use our data. Thus, as a robustness check, we experiment with an alternative specification in which we aggregate the information at the company-state-occupation-month level and calculate the share of remote work vacancies and the share of remote work vacancies that provide performance performance pay. We then regress the latter on the former, so that the correlation is identified only through within company-occupation variation over time. The results are presented in Online Appendix Table B1, which shows that both the 2SLS estimates are positive and significant at the 99 percent confidence level. The coefficient implies that on average, one quarter of the new remote vacancies provide performance pay.

6 Conclusions

The surge in remote work has introduced new information asymmetries, with a reduced number of workers physically present in the office, thereby impeding managers' ability to effectively monitor their actions. This paper investigates how firms have addressed this challenge, with a particular focus on the utilization of variable compensation as a means to align incentives. Drawing on the Principal-Agent model with moral hazard, which predicts that remote work should increase firms' reliance on performance-based pay, we empirically explore this relationship using dictionary methods and an extensive dataset comprising nearly all online job postings in the United States between 2018 and 2022.

Doing so, we find that remote vacancies are twice as likely to provide performance pay, relative to onsite ones. This finding holds noteworthy implications, not least since hampering the insurance-providing role of companies might have far-reaching economic consequences for remote workers. In particular, because employees cannot easily insure against idiosyncratic shocks through market mechanisms due to moral hazard, this might impact their access to credit, portfolio allocation, and labor supply.

References

Adrjan, P., Ciminelli, G., Judes, A., Koelle, M., Schwellnus, C. & Sinclair, T. (2021), 'Will it stay or will it go? analysing developments in telework during covid-19 using online job postings data'.

Aggarwal, R. K. & Samwick, A. A. (1999), 'The other side of the trade-off: The impact of risk on executive compensation', *Journal of political economy* **107**(1), 65–105.

Aksoy, C. G., Barrero, J. M., Bloom, N., Davis, S. J., Dolls, M. & Zarate, P. (2022), Working from home around the world, Technical report, National Bureau of Economic Research.

Altig, D., Baker, S., Barrero, J. M., Bloom, N., Bunn, P., Chen, S., Davis, S. J., Leather, J., Meyer, B., Mihaylov, E. et al. (2020), 'Economic uncertainty before and during the covid-19 pandemic', *Journal of Public Economics* **191**, 104274.

Autor, D., Dube, A. & McGrew, A. (2023), The unexpected compression: Competition at work in the low wage labor market, Technical report, National Bureau of Economic Research.

Azariadis, C. (1975), 'Implicit contracts and underemployment equilibria', *Journal of political economy* **83**(6), 1183–1202.

Baily, M. N. (1974), 'Wages and employment under uncertain demand', *The Review of Economic Studies* **41**(1), 37–50.

Baker, S. R., Bloom, N., Davis, S. J. & Terry, S. J. (2020), Covid-induced economic uncertainty, Technical report, National Bureau of Economic Research.

Barrero, J. M., Bloom, N. & Davis, S. J. (2021), Why working from home will stick, Technical report, National Bureau of Economic Research.

Barrero, J. M., Bloom, N., Davis, S. J., Meyer, B. H. & Mihaylov, E. (2022), The shift to remote work lessens wage-growth pressures, Technical report, National Bureau of Economic Research.

Bloom, N., Han, R. & Liang, J. (2022), How hybrid working from home works out, Technical report, National Bureau of Economic Research.

Bloom, N., Liang, J., Roberts, J. & Ying, Z. J. (2015), 'Does working from home work? evidence from a chinese experiment', *The Quarterly Journal of Economics* **130**(1), 165–218.

Bojinov, I., Choudhury, P. & N Lane, J. (2021), 'Virtual watercoolers: A field experiment on virtual synchronous interactions and performance of organizational newcomers', *Harvard Business School Technology & Operations Mgt. Unit Working Paper* (21-125).

Brinatti, A., Cavallo, A., Cravino, J. & Drenik, A. (2021), The international price of remote work, Technical report, National Bureau of Economic Research.

Brynjolfsson, E., Horton, J. J., Ozimek, A., Rock, D., Sharma, G. & TuYe, H.-Y. (2020), Covid-19 and remote work: An early look at us data, Technical report, National Bureau of Economic Research.

Choudhury, P., Foroughi, C. & Larson, B. (2021), 'Work-from-anywhere: The productivity effects of geographic flexibility', *Strategic Management Journal* **42**(4), 655–683.

Choudhury, P., Khanna, T., Makridis, C. & Schirmann, K. (2022), 'Is hybrid work the best of both worlds? evidence from a field experiment', *Harvard Business School Technology & Operations Mgt. Unit Working Paper* (22-063), 22–063.

Criscuolo, C., Gal, P., Leidecker, T., Losma, F. & Nicoletti, G. (2021), 'The role of telework for productivity during and post-covid-19: Results from an OECD survey among managers and workers', *Working Paper* .

Dalton, M., Dey, M. & Loewenstein, M. (2022), 'The impact of remote work on local employment, business relocation, and local home costs', *Working Paper 553* .

Guiso, L., Pistaferri, L. & Schivardi, F. (2005), 'Insurance within the firm', *Journal of Political Economy* **113**(5), 1054–1087.

Hansen, S., Lambert, P. J., Bloom, N., Davis, S. J., Sadun, R. & Taska, B. (2023), Remote work across jobs, companies, and space, Technical report, National Bureau of Economic Research.

Holmström, B. (1979), 'Moral hazard and observability', *The Bell journal of economics* pp. 74–91.

Holmström, B. (2017), 'Pay for performance and beyond', *American Economic Review* **107**(7), 1753–1777.

Holmstrom, B. & Milgrom, P. (1987), 'Aggregation and linearity in the provision of intertemporal incentives', *Econometrica: Journal of the Econometric Society* pp. 303–328.

Holmstrom, B. & Milgrom, P. (1991), 'Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design', *The Journal of Law, Economics, and Organization* **7**(special_issue), 24–52.

Jensen, M. C. & Murphy, K. J. (1990), 'Performance pay and top-management incentives', *Journal of political economy* **98**(2), 225–264.

Juhn, C., McCue, K., Monti, H. & Pierce, B. (2018), 'Firm performance and the volatility of worker earnings', *Journal of Labor Economics* **36**(S1), S99–S131.

Kwan, A. & Matthies, B. (2022), 'Measuring the impact of remote work using big data', *Working Paper* .

Lazear, E. P. (1986), 'Salaries and piece rates', *Journal of business* pp. 405–431.

Lazear, E. P. (2000), 'Performance pay and productivity', *American Economic Review* **90**(5), 1346–1361.

Milgrom, P. R. (1981), 'Good news and bad news: Representation theorems and applications', *The Bell Journal of Economics* pp. 380–391.

Mirrlees, J. A. (1999), 'The theory of moral hazard and unobservable behaviour: Part i', *The Review of Economic Studies* **66**(1), 3–21.

Rosen, S. (1985), 'Implicit contracts: a survey'.

Online Appendix

A Model Appendix

A.1 Working at The Office: Full Information

With full information, the company determines the level of effort that maximizes her utility (net profits) and the compensation scheme that will induce exactly that level of effort. Contract theory suggests that with full information the company can design a "forcing contract" specifying the desired effort e^* and threaten the employee by setting $s(x,e) = -\infty$ if $e \neq e^*$. In this case, the only constraint of the company is the participation constraint of the employee. Therefore, the company's optimization problem is

$$\max_{s(x)} \quad E\big[x - s(x)|e^*\big]$$

subject to

$$E\big[U\big(s(x)\big)|e^*\big] \geq \underline{\mathrm{H}} + c(e^*)$$

Letting λ be the Lagrangean multiplier associated to this problem and differentiating the Lagrangean with respect to $s(x)$, yields the optimality condition

$$\frac{1}{U'\big(s(x)\big)} = \lambda \tag{A1}$$

Since $\lambda > 0$ is constant, we see that under full information, the optimal contract is a fixed payment conditional on exerting the optimal level of effort.[19]

[19] At the optimum, the participation constraints of the employee is binding, which implies that $\lambda > 0$. The binding participation constraint gives the optimal fixed compensation $s = {U'}^{-1}\big(\underline{\mathrm{H}} + c(e^*)\big)$.

A.2 Working from Home: Hidden Action

Unlike in the case of full information, the set of observable and contractible events Ω includes x, but no longer e. In other words, the optimal contract cannot be conditioned on e as the latter is not observed by the company.

We follow Holmström (2017) in thinking about the problem in terms of the conditional distribution of x given e, $F(x|e)$, which is induced on x by the variation in θ. We assume that F is twice continuously differentiable with density function f, and that its derivative with respect to effort f_e exists.

The problem of the company becomes

$$\max_{s(x),e} \int \big[x - s(x)\big]\, dF(x|e)$$

subject to

$$\int \big[U\big(s(x)\big) - c(e)\big]\, dF(x|e) \geq \underline{\mathrm{H}}$$

$$\int U\big(s(x)\big) f_e(x|e)\, dx = c'(e) \tag{A2}$$

Condition (A2) in the new problem is the incentive compatibility constraint (ICC) of the employee. Two remarks are warranted. First, the ICC is now needed because the company cannot observe employees' effort, which implies that she has to rely on their selfish behavior in order to achieve her desired level. Second, to specify the ICC, we use the first order conditions of the employee with respect to his optimal effort level. Therefore, we are implicitly assuming that a solution of that problem exists and it is unique.[20]

Letting λ be the Lagrangean multiplier associated to the participation constraint, μ the multiplier associated to the ICC, and differentiating the Lagrangean with respect to $s(x)$,

[20] In the literature, this is referred to as the "first order approach". See for a discussion of its limitations.

yields the optimality condition

$$\frac{1}{U'\big(s(x)\big)} = \lambda + \mu \cdot g(x|e) \tag{A3}$$

where $g(x|e) \equiv f_e(x|e)/f(x|e)$. Condition (A1) differs from (A3) in that it has a variable term that depends on performance x.[21] Therefore, if working from home introduces a problem of hidden action because employees cannot be monitored as well as when they work at the office, we expect companies to rely more on performance-based compensation.

A.2.1 Uncertainty

Milgrom (1981) shows that when $g(x|e)$ in Equation A3 is monotonically increasing in x, then $\mu > 0$. This implies that $s(x)$ is also increasing in measured performance, the larger the deviation from a given level of effort.[22] Thus, employees are rewarded for higher-than-expected realizations of x because these are interpreted as evidence of higher-than-expected levels of effort, and vice versa.

If we assume that $F(x|e)$ is a normal distribution with mean e and variance σ^2, then

$$g(x|e) = \frac{x - e}{\sigma^2} \tag{A4}$$

Equation (A4) says that the variable component of the optimal contract depends inversely on the variance of the measurement error. The lower the precision with which the company can observe employees' performance, the less room for performance pay compensation.

[21] As in (A1), $\lambda > 0$ due to the participation constraint being binding. Thus, $\mu = 0$ would imply a fixed payment that would induce the employee to exert the lowest possible level of effort under hidden action, which cannot be optimal except than in the uninteresting case in which $e = 0$ is optimal.

[22] Marginal Likelihood Ratio Property (MLRP).

B Figures and Tables Appendix

Figure B1: Share of performance pay and Pandemic Uncertainty Index.

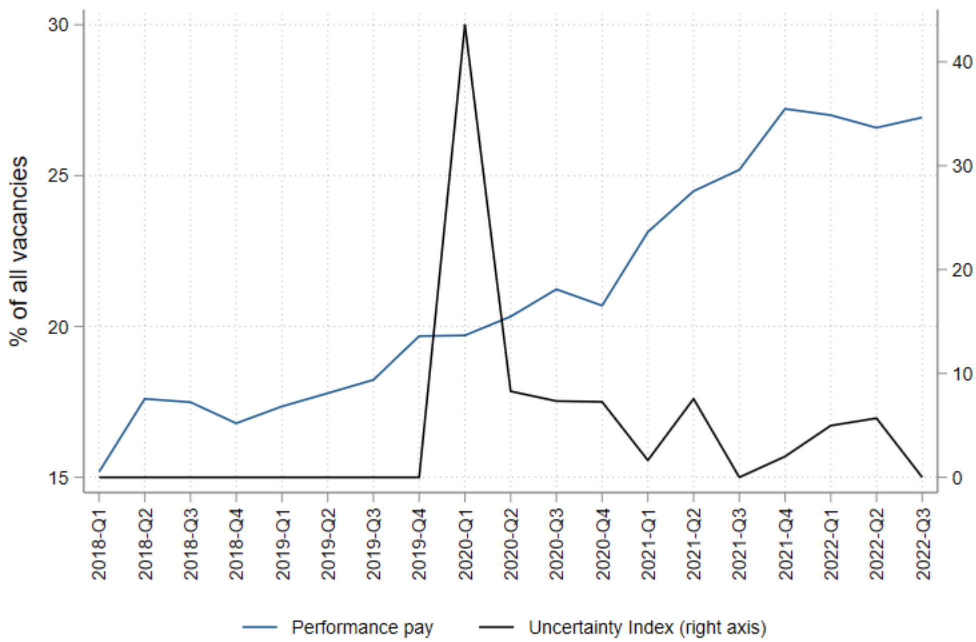

Notes: This figure presents the sample average share of performance pay online job postings by broad occupational category. To tag performance pay vacancies, we use dictionary methods based on the keywords in Table 2. The Pandemic Uncertainty Index is taken from Baker et al. (2020).

Figure B2: Average difference in share of remote job postings by recruiters in mandating *vs* non-mandating states.

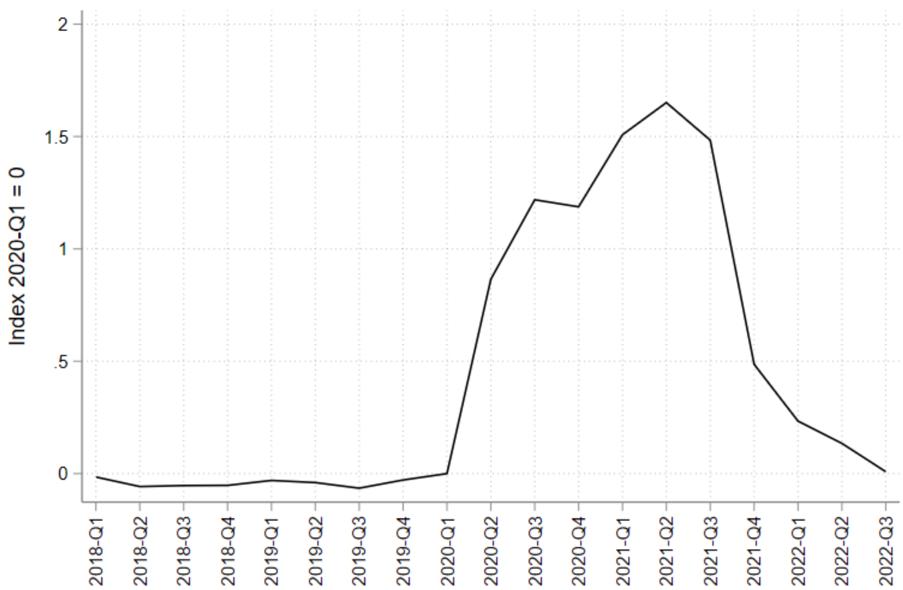

Notes: This figure presents the share of remote work vacancies over quarters by mandating and non-mandating states of strict workplace closing rules.

Figure B3: Workplace closing and share of remote jobs (red dashed line).

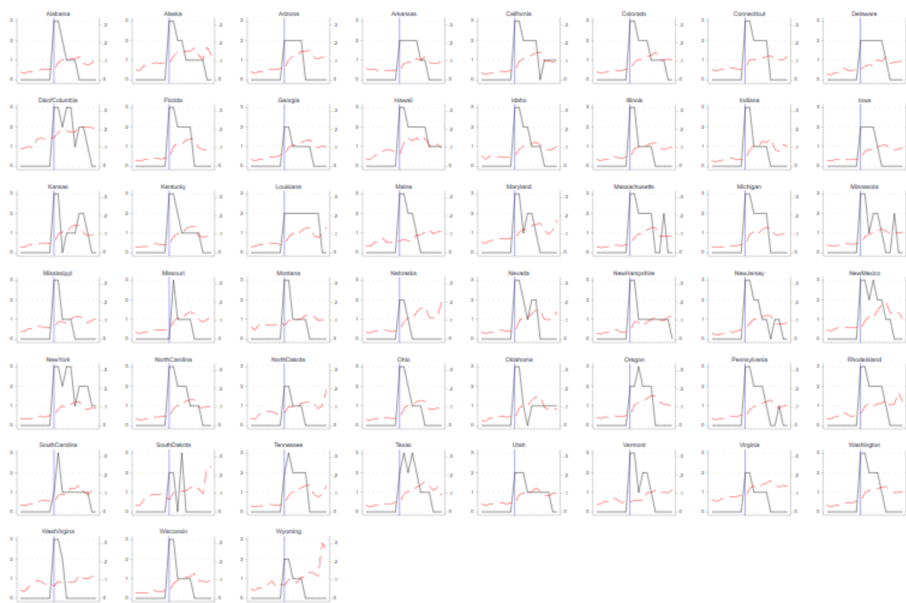

Notes: This figure presents the evolution of workplace closing rules and the share of remote vacancies by US state.

Table B1: Company-occupation-month aggregation.

	(1) Performance pay (share)	(2) Remote (share)	(3) Performance pay (share)
(mean) remote	0.256***		0.276***
	(0.000)		(0.018)
ivclosed		0.006***	
		(0.000)	
Observations	44,946,724	44,946,724	44,946,724
R-squared	0.210	0.556	
Company-occupation FE	yes	yes	yes
Month FE	yes	yes	yes
State-level controls	no	yes	yes
First stage F			777.2

Notes: This table presents 2SLS estimates of system (3)-(4). The instrument $\Psi_s = 1$ if state s mandates closing in March 2020, and zero otherwise. The vector X_{st} includes log-real state-level GDP and the log of COVID-19 cases. To tag remote job vacancies, we use dictionary methods based on the keywords in Table 1. To tag performance pay vacancies, we use dictionary methods based on the keywords in Table 2. Standard errors are clustered at the company-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table B2: Fixed effects including 6-digit NAICS industry.

	(1) Remote job	(2) Performance pay
Remote job		0.311***
		(0.084)
State mandates workplace closing (yes = 1)	0.003***	
	(0.000)	
Observations	31,845,177	31,845,177
R-squared	0.323	
Company-occupation FE	yes	yes
Month FE	yes	yes
State-level controls	yes	yes
First stage F		304

Notes: This table presents 2SLS estimates of system (3)-(4). The instrument $\Psi_s = 1$ if state s mandates closing in March 2020, and zero otherwise. The vector X_{st} includes log-real state-level GDP and the log of COVID-19 cases. To tag remote job vacancies, we use dictionary methods based on the keywords in Table 1. To tag performance pay vacancies, we use dictionary methods based on the keywords in Table 2. Standard errors are clustered at the company-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table B3: Full sample including laggard states and all time periods.

	(1) Remote job	(2) Performance pay
Remote		0.256***
		(0.030)
State mandates workplace closing (yes = 1)	0.006***	
	(0.000)	
Observations	108,475,042	108,475,042
R-squared	0.496	
Company-occupation FE	yes	yes
Month FE	yes	yes
State-level controls	yes	yes
First stage F		2172

Notes: This table presents 2SLS estimates of system (3)-(4). The instrument $\Psi_s = 1$ if state s mandates closing in March 2020, and zero otherwise. The vector X_{st} includes log-real state-level GDP and the log of COVID-19 cases. To tag remote job vacancies, we use dictionary methods based on the keywords in Table 1. To tag performance pay vacancies, we use dictionary methods based on the keywords in Table 2. Standard errors are clustered at the company-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.