

Who Should Work from Home during a Pandemic? The Wage-Infection Trade-off

Sangmin Aum^{*} Sang Yoon (Tim) Lee[†] Yongseok Shin[‡]

October 6, 2020

Abstract

Shutting down the workplace is an effective means of reducing contagion, but can incur large economic losses. We construct an exposure index, which measures infection risks across occupations, and a work-from-home index, which gauges the ease with which a job can be performed remotely across both industries and occupations. Because the two indices are negatively correlated but distinct, the economic costs of containing a pandemic can be minimized by only sending home those jobs that are highly exposed but easy to perform from home. Compared to a lockdown of all non-essential jobs, the optimal policy attains the same reduction in aggregate exposure (32 percent) with one-third fewer workers sent home (24 vs. 36 percent) and with only half the loss in aggregate wages (15 vs. 30 percent). A move from the lockdown to the optimal policy reduces the exposure of low-wage workers the most and the wage loss of the high-wage workers the most, although everyone's wage losses become smaller. A constrained optimal policy under which health workers cannot be sent home still achieves the same exposure reduction with a one-third smaller loss in aggregate wages (19 vs. 30 percent).

JEL Classifications: *E24, I14, J21*

Keywords: Lockdown, exposure, work from home, pandemic

^{*}Myongji University; aumsang@gmail.com

[†]Queen Mary University of London, TSE and CEPR; sylee.tim@qmul.ac.uk. Lee acknowledges financial support from the British Academy [grant number COV19\201483].

[‡]Washington University in St. Louis, Federal Reserve Bank of St. Louis and NBER; yshin@wustl.edu

1 Introduction

Lockdowns and social distancing have become the most discussed policy tools since the emergence of COVID-19. Most countries have implemented such measures to varying degrees to contain the pandemic. The obvious downside is their economic costs, since most economic activities depend on in-person interactions. Thus at least in the short run, policymakers face an inherent trade-off between the risk of contagion and economic losses.

To analyze this trade-off, it is important to know, first, the exposure or the infection risk of performing a given job and, second, the ease with which the job can be performed remotely. The actual trade-off will depend on how jobs are distributed along these two dimensions, which is the focus of this paper.

We begin by constructing an index of exposure to infection risks across occupations from the O*NET, and an index of time spent working from home (WFH) across industries and occupations from the American Time Use Survey (ATUS). Such indices are not new, but our approach is novel in that we cross-walk and merge the two datasets in order to quantify the trade-off between the infection risks and economic losses in the economy, as represented by the distribution of workers in the American Community Survey (ACS).

We find that, although jobs with a low WFH tend to be more exposed to infection on average, the correlation of the two across occupations and industries is far from tight. Infection risks vary widely even among jobs with the same WFH: for example, neither medical therapists nor experimental physicists can work from home, but the latter pose almost no risk of contagion. Conversely, there are a number of jobs with high infection risk that can be effectively performed from home, such as an IT sales agent. In addition, even the same occupation can have a very different WFH index depending on the industry: for example, a registered nurse employed by a hospital has a low WFH, but one in consulting services has a high WFH.

Compared to a blanket closure of broad categories of jobs (essential vs. non-essential, in particular), then, targeting social distancing toward workers in specific occupations of detailed industries can substantially reduce the economic cost, while still attaining the same reduction in the aggregate exposure to infection. To quantify such possible gains, we consider an optimal policy whereby the planner chooses which industry-occupation pairs to send home in order to minimize the aggregate wage loss, subject to a given reduction in the aggregate exposure

to infection. Intuitively, it is optimal to first send home jobs with high exposure at work and low wage loss when working from home, the latter of which can be computed from the wage and WFH of a job. Mathematically, this translates to the optimal policy being characterized by a linear threshold in the 2-dimensional plane of wage loss and exposure.

The aggregate wage loss under the optimal policy is much smaller than under a lockdown of all non-essential jobs as implemented in many U.S. states and European economies. Our version of the real-world lockdown reduces aggregate exposure by 32 percent by sending home 36 percent of all workers, costing 30 percent of aggregate wages. Our optimal policy attains the same reduction in aggregate exposure by sending home only 24 percent of all workers, costing only 15 percent of aggregate wages. That is, the optimal policy achieves the same reduction in aggregate infection risk at half the economic cost, while a third fewer workers are sent home. Under a constrained optimal policy in which health workers must continue to work normally, the aggregate wage loss is 19 percent, still a third less than under a real-world lockdown. These gains are possible because of the large variation of WFH across occupations and industries, for any given level of exposure.

It has become clear that low-wage workers are not only bearing the brunt of the pandemic economically, but also in terms of infection risks. Compared to the lockdown of all non-essential jobs, the optimal policy reduces low-wage workers' infection risks the most. On the other hand, a move from the lockdown to the optimal policy generates the largest wage gains for high-wage workers, although more workers across the entire wage distribution are allowed to work normally and thus earn more. In fact, under all policy scenarios we consider, it is always high-wage workers who are the least exposed to infection risks and also lose the least economically, pointing to the importance of redistributive policies during a pandemic.

A final word of caution: By design, our optimal policy is simple and static, abstracting from the essentiality of certain jobs (other than healthcare providers) that need to be performed even in the midst of a pandemic, the complementary among jobs that need to be performed in-person, the economic propagation across jobs and sectors, and the possibility that people switch to jobs with lower exposure or lower work-from-home wage losses. It also assumes that the indices we construct are constant, ignoring the potential change in exposure or WFH of specific jobs due to more subtle non-pharmaceutical interventions (NPI) such as

wearing masks or reorganizing the workplace or changes in individual behaviors as the population adjusts to the pandemic. Nevertheless, our analysis presents a simple guidance for policymakers that is easy to implement in practice, while also providing a benchmark for structural economic models that consider some of these dimensions that we abstract from.¹

2 Data

2.1 Employment and Wage by Occupation and Industry

We compute employment weights and mean hourly wages by occupation and industry from the ACS. Following [Acemoglu and Autor \(2011\)](#), we only include civilian, prime age workers (between 16 and 65 years old). An individual’s employment weight is their sampling weight multiplied by their usual annual hours worked (usual hours worked in a week times usual weeks worked in a year). For each year, we multiply top-coded wages by 1.5, and bottom-code the lowest hourly wage percentiles. Then the hours-adjusted employment weight and mean hourly wage of each occupation-industry combination are computed using consistent industry and occupation codes following [Autor and Dorn \(2013\)](#), modified to incorporate changes in the Census industry (IND) and occupation codes (OCC) from 2014 to 2018. Finally, we take a simple average of the employment weights and hourly wages over the 5 years.

2.2 Constructing Exposure and Work-from-Home Indices

Exposure Index The O*NET asks experts and workers to give numerical answers to questions that capture detailed characteristics of an occupation, where an occupation is defined by its Standard Occupation Classification (SOC) code. To construct our exposure index, we take the weighted average of the answers to 2 questions: “Physical Proximity” (PP) and “Exposed to Disease or Infections” (EDI).²

We first convert O*NET titles to SOC codes using their accompanying cross-

¹A review of the relevant literature is in the appendix.

²These are questions 4.C.2.a.3 and 4.C.2.c.1.b, respectively. These two descriptors were analyzed in detail by the New York Times ([Gamio, 2020](#)) for the U.S. and reproduced for the U.K. by its Office for National Statistics (ONS). The New York Times and ONS focused on each measure separately and showed that the two are strongly positively correlated. The economics literature has only used physical proximity as a measure of exposure ([Leibovici et al., 2020](#); [Mongey et al., 2020](#)).

walk.³ The SOC codes are then mapped to ACS OCC codes using a crosswalk available from the IPUMS USA, which is heavily modified so that each ACS OCC code has a unique value for both descriptors.⁴ Finally, we normalize PP and EDI to have a mean of 0 and standard deviation of 1, and take the average of the two as our exposure index.

Work-from-Home Index Some earlier studies on COVID-19 used O*NET descriptors to construct a WFH index (e.g. Dingel and Neiman, 2020; Mongey et al., 2020). However, a better measure of the ease with which a job can be performed from home may be whether people actually do the job from home, such as reported in ATUS (Adams-Prassl et al., 2020). While Mongey et al. (2020) do show that the former is positively correlated with the latter, there are both qualitative and quantitative reasons to favor the latter. Qualitatively, some descriptors included in O*NET-based WFH indices are misleading: For example, O*NET’s “Outdoor” categories (implying low WFH) include farmers who in fact have a high incidence of actually working from home, since they work in self-owned plots and land. Quantitatively, ATUS allows WFH to vary across industries, as well as across occupations, and indeed even the same occupation has very different WFH across industries in the data.⁵

For each industry-by-occupation (IND-OCC) combination in ATUS, we compute the total time spent working and total time worked from home across all individuals in the corresponding cell for each year from 2014 to 2018, using each year’s sampling weight. We then take a simple average over all 5 years, and the ratio of the average time worked from home to the average time spent working is our WFH index. Each cell is then matched to the ACS, and in this process we merge and impute missing cells using ACS employment weights, and also ensure that each OCC code has non-missing O*NET descriptors.

All in all, we end up with 458 occupations with a unique exposure index, spread over 254 industries. Each of the 458×254 OCC-IND cells has a unique WFH index, as well as total hours worked and hourly wages.

³Available at <https://www.onetcenter.org/crosswalks.html>.

⁴In general, O*NET SOC is finer than OCC, so each descriptor for a lower level SOC occupation is averaged and subsumed into a higher level SOC occupation using the SOC-OCC crosswalk. However, the O*NET does not list descriptors for any lower level of some of the 3-digit OCC codes, and OCC codes changed in 2018, necessitating additional manipulations.

⁵Adams-Prassl et al. (2020) show that this is also true in the United Kingdom.

2.3 Relationship between Exposure and WFH

Are jobs with higher exposure to infection risks harder to perform from home? In Figure 1, each circle represents a job, defined as a specific occupation in a specific industry, and its location shows its WFH index (horizontal axis) and exposure index (vertical axis). The size of the circle denotes the (hours-weighted) employment share of each job, averaged from 2014 to 2018.

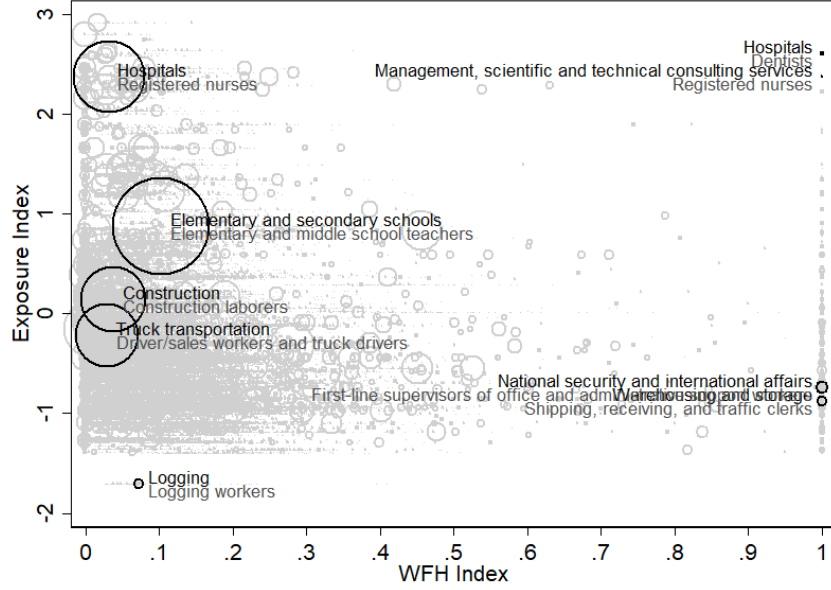
Three patterns emerge. First, consistent with conventional wisdom, more exposed jobs tend not to be performed from home. For example, hospital nurses are highly exposed to infection, and at the same time, rarely work from home. Second, the negative correlation is not very tight. In Figure 1, there are many industry-occupation pairs *not* working from home, regardless of their exposure to infection risks. Most notably, loggers are one of the least exposed occupations, but for obvious reasons do not work from home. Third, even the same occupation shows substantial variations in WFH across industries. For example, while hospital nurses do not work from home, a small number of registered nurses in the consulting service industry exclusively work from home.

The table in Figure 1 shows the coefficients from regressing the WFH index on the exposure index for our whole industry \times occupation sample (first column), as well as only by occupation (second column) and only by industry (third column). For the latter cases, industry average exposure or WFH indices are computed by averaging across each occupation’s exposure or WFH within an industry, weighted by each occupation’s within-industry employment shares, and similarly for the indices by occupation. In all three cases, the negative correlation between the exposure and WFH indices is statistically significant. However, a large fraction of the dispersion in one index still remains unexplained by the other, as represented by the low R^2 ’s in all regressions. In particular, the R^2 is as low as 0.03 for the full sample of industry-occupation combinations. Since the exposure index varies only across occupations but not industries by construction, the large reduction in R^2 from the second to first column mirrors the wide variation in WFH across industries even for the same occupation.⁶

3 Optimal Policy

We interpret the fact that workers in a job are more likely to work from home (high WFH) to mean that the productivity loss from working from home is

⁶Appendix C shows that the correlation between our indices and demographics are also weak.



| | Work-from-Home Index | | |
|----------|----------------------|-------------------|-------------------|
| | Ind. by Occ. | By Occupation | By Industry |
| Exposure | -0.033 (0.004) | -0.033 (0.005) | -0.045 (0.010) |
| Obs. | 54,108 | 458 | 254 |
| R^2 | 0.034 | 0.109 | 0.126 |

Fig. 1: Relationship between Exposure and WFH

- 1) Each circle represents a specific occupation within a specific industry (458 occupations \times 254 industries). The size of the circles denotes the hours-adjusted employment share, averaged from 2014 to 2018. Dark blue circles are some examples either with very large employment shares or with extreme values of exposure and/or WFH. The first description for these examples is the industry, and the second the occupation.
- 2) The table shows the result of regressing WFH on exposure, with observations weighted by their hours-adjusted employment share. Robust standard errors are in parentheses. The unit of observation is an industry-occupation pair in the first column, occupation in the second column, and industry in the last column.

smaller for the said job. Thus, the large dispersion in WFH among jobs with similar levels of exposure implies that the economic cost of sending home workers to reduce infection risks will depend on whether they work in a high or low WFH job.

Let (s_i, w_i, e_i, h_i) denote the employment share, average wage, exposure index and WFH index for each industry-occupation combination i . An optimal policy minimizes the economic cost from sending home a fraction $0 \leq x_i \leq 1$ of each industry-occupation combination, subject to reducing aggregate exposure by at least a given fraction $0 \leq y \leq \psi$:

$$\min_{\{x_i \in [0,1]\}_{i=1}^I} \psi \sum_{i=1}^I s_i(1 - h_i)w_i x_i \quad \text{s.t.} \quad \psi \sum_{i=1}^I s_i e_i x_i \geq y \sum_{i=1}^I s_i e_i, \quad (1)$$

where ψ is a “compliance” constant that measures how many workers in the jobs being sent home do comply with the policy and work from home, and I is the total number of occupation-industry combinations.⁷ This problem is merely a high-dimensional linear programming problem.⁸ For the problem to be well-defined, all e_i ’s need to be non-negative, so we simply shift our exposure index by subtracting off its minimum value. The problem implies that the economic cost of working from home is proportional to $(1 - h_i)w_i$, where h_i is our WFH index that lies between 0 and 1. Since the problem is linear, the solution x_i^* will be either 0 or 1 for all i , unless the job happens to fall exactly on the threshold.

Figure 2 shows the optimal policy when $\psi = 0.487$.⁹ In panels (a) and (b), each industry-occupation combination is plotted as a circle along the exposure (e_i , vertical axis) and the wage loss from working remotely ($w_i(1 - h_i)$, horizontal axis) dimensions. The size of a circle represents job i ’s employment share. The optimal policy is a linear threshold in this two-dimensional space, where only jobs that are above the threshold are sent home (blue circles). The slope of the threshold is positive since the optimal policy takes into account both the exposure and the wage loss from working remotely. That is, even if a job has high infection risks, it is not sent home if it is hard to be done remotely (low WFH) and has a high average wage, implying a large wage loss from working remotely. Panel (a)

⁷Thus the maximal possible reduction in aggregate exposure is ψ . This may be due to non-compliance, selective furloughs, or reduction in working hours (such as a curfew on pubs and restaurants) as opposed to a shutdown order, and so on. Note that the value of ψ does not affect the optimal solution $\{x_i^*\}$ but only affects the magnitude of the wage loss and exposure reduction.

⁸Program (1) is the dual problem of minimizing exposure subject to a given level of aggregate wage losses.

⁹ $\psi = 0.487$ corresponds to a 30 percent drop in aggregate wages under the real-world lockdown, to be shown in Section 4.

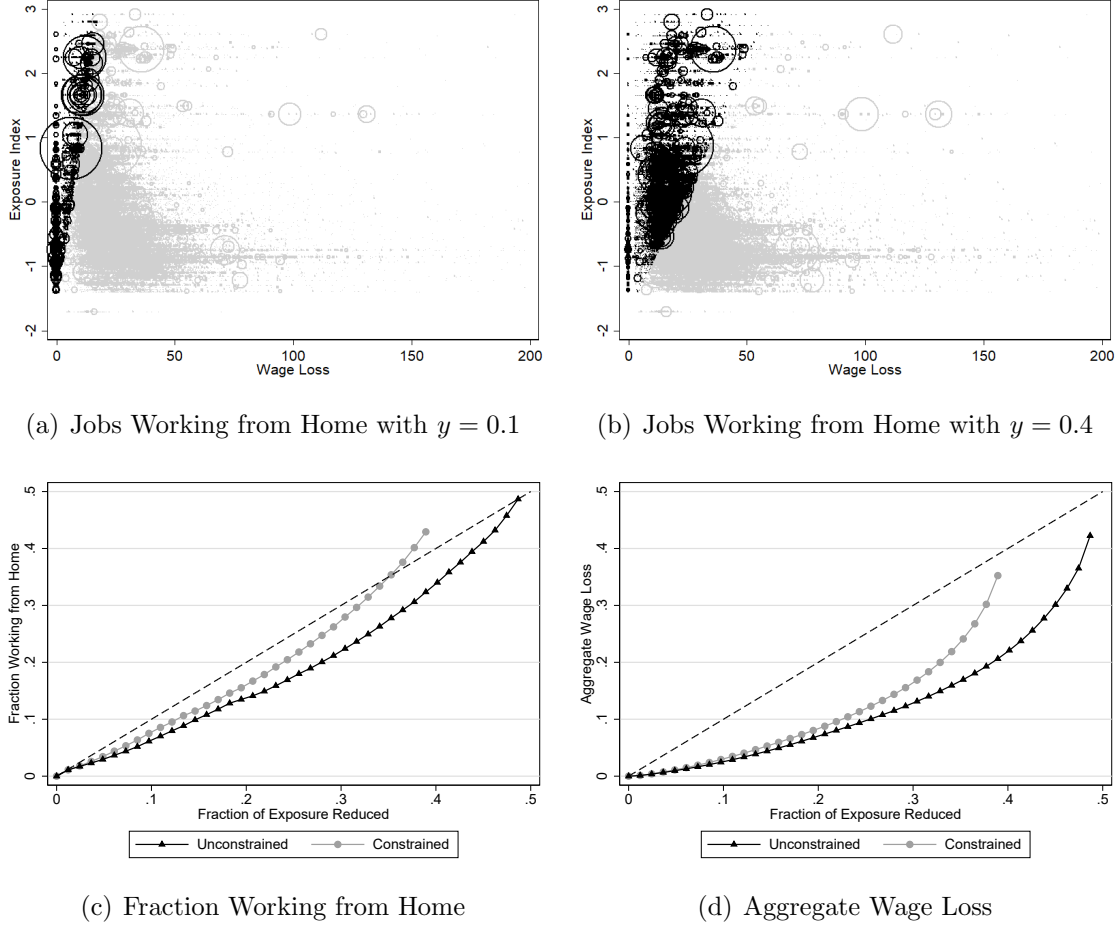


Fig. 2: Optimal Policy

In panels (a) and (b), each circle is a specific occupation within a specific industry (458 occupations \times 254 industries). The size of the circles denotes the hours-adjusted employment share, average from 2014 to 2018. The blue circles represent jobs that are sent home by the optimal policy (group i such that $x_i^* = 1$) and the rest are performed normally. Panel (a) is when the aggregate exposure is reduced by 10 percent, and (b) is by 40 percent. In the lower panels, for a fraction y of reduction in population exposure on the horizontal axis, the fraction of workers sent home and the aggregate wage loss (with the scaling constant $\psi = 0.487$ under the optimal policy) are shown as blue lines in panels (c) and (d), respectively. The constrained optimal policy solves the same problem (1), but with the added constraint that all health workers must work normally, and are shown as red lines.

is the solution when aggregate exposure is reduced by 10 percent, and panel (b) by 40 percent.

Panel (c) plots the fraction of workers working from home under the optimal policy in blue, for each level of reduction in aggregate exposure (y) on the horizontal axis. Panel (d) shows the aggregate wage loss from the optimal policy in blue against the same horizontal axis. Both blue lines are upward sloping, since reducing exposure requires sending more workers home, which also leads to larger wage losses. More important, the lines are convex, because the optimal policy sends home those jobs with higher exposure and lower wage losses first.

One unpleasant feature of the optimal policy we solve above is that many health workers are sent home because they tend to have very high exposure, even though most of them cannot work from home and thus incur large wage losses if compelled to do so. For a more realistic problem during a pandemic, we also solve a constrained optimal policy in which health workers must work normally. These are workers in the healthcare and social assistance sectors (IND codes 7970-8390), and healthcare practitioners, technical, and support occupations (OCC codes 3000-3655), who comprise about 11 percent of (hours-adjusted) aggregate employment. Combined, these jobs account for about 20 percent of the aggregate exposure.

In the two bottom panels of Figure 2, the red lines plot the fraction of workers working from home under the constrained optimal policy in panel (c) and the aggregate wage loss in panel (d), both against the reduction in aggregate exposure y on the horizontal axis. Compared to the unconstrained optimal policy in blue, the constrained policy sends more workers home and causes more wage losses for any given reduction in aggregate exposure, but is still visibly convex.

4 Optimal Policy vs. Real-world Lockdowns

We now compare our optimal policy with a lockdown that mimics those implemented in the U.S. and most European countries. Although lockdowns were implemented with varying degrees of severity across countries and U.S. states, they did share common features. Most often, the government classified industries (and/or occupations) as essential or non-essential, and tried to keep essential workers working as normally as possible, while forbidding non-essential workers from commuting to work. For example, the Cybersecurity and Infrastructure Agency (CISA) of the U.S. Department of Homeland Security provides guide-

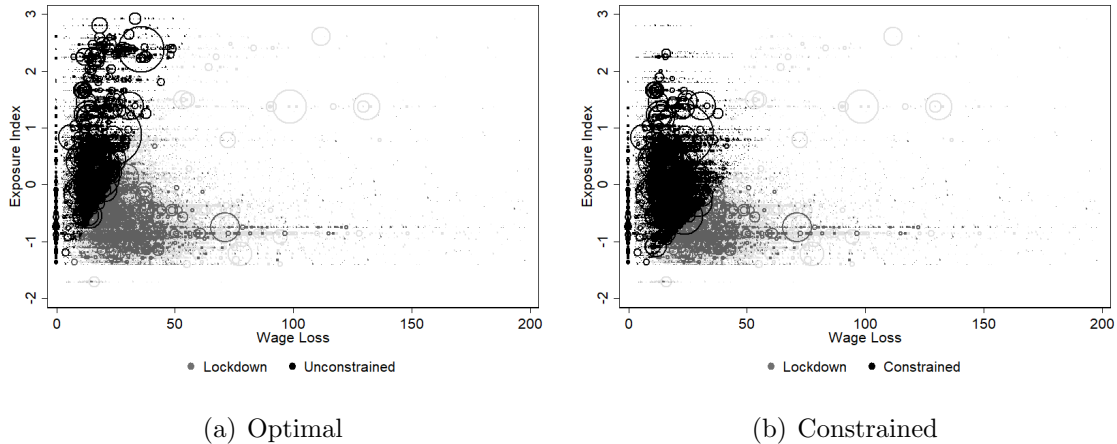


Fig. 3: Optimal Policy vs. Lockdown

Each circle is a specific occupation within a specific industry (458 occupations \times 254 industries). The size of a circle represents its employment share. Red circles indicate jobs sent home during the actual lockdown. Blue circles are the jobs done remotely under the unconstrained optimal policy in the left panel and under the constrained optimal policy in the right panel.

lines on which jobs are essential for critical infrastructure.

Our version of the real-world lockdown follows [Palomino et al. \(2020\)](#), who show which occupations and industries were effectively locked down in Europe.¹⁰ In the context of our problem (1), locked down jobs are a set $\{i \in \{1, 2, \dots, I\} | x_i = 1\}$, which in general will differ from the optimal solution. We set $\psi = 0.487$ as we did in the previous section, which generates a 30-percent drop in aggregate wages under the actual lockdown.¹¹ Given ψ , we find that the lockdown reduces aggregate exposure by 32 percent.

Panel (a) of Figure 3 shows which jobs were non-essential and hence ordered to stay home during the actual lockdown (red circles) and which jobs are performed from home under the optimal policy (blue circles). Both achieve the same reduction in aggregate exposure, but there is not much overlap between the two policies in terms of which jobs are sent home. Panel (b) shows the jobs sent home under the constrained optimal policy (i.e., all health workers working normally), instead of the optimal policy.

Table 1 shows exactly how this difference manifests in terms of the fraction of workers sent home and the associated aggregate wage loss. The gains from

¹⁰[Palomino et al. \(2020\)](#) identify essential jobs by ISCO (2-digit) and NACE (1-digit) from the lockdowns implemented in Italy and Spain. We match these jobs to the OCC and IND codes in the Census. Their list of essential jobs are broadly consistent with the CISA guidelines.

¹¹With a labor income share of 60 percent, this implies an 18-percent drop in GDP due to the lockdown, which is between the GDP loss in the U.S. (10 percent) and that in the U.K (20 percent) from the first to second quarter of 2020.

| | Fraction Working Remotely | Wage Loss |
|-------------------------|---------------------------|-----------|
| (A) Lockdown | 0.359 | 0.300 |
| (B) Optimal | 0.243 | 0.145 |
| (C) Constrained optimal | 0.306 | 0.191 |

Table 1: Policy Outcomes

The first column is the (hours-adjusted) fraction of workers working from home, and the second the aggregate wage loss.

implementing the optimal policy are substantial: The same reduction in exposure can be attained by sending home one-third fewer workers (36 percent working from home by the lockdown vs. 24 percent under the optimal policy) at half the economic cost (aggregate wage loss of 30 percent by the lockdown vs. 15 percent under the optimal policy).

Even with the constraint that all health workers continue to work normally, the constrained optimal policy does substantially better than the actual lockdown. It delivers the same reduction in aggregate exposure with a one-third smaller loss in aggregate wages (19 percent instead of 30 percent).¹²

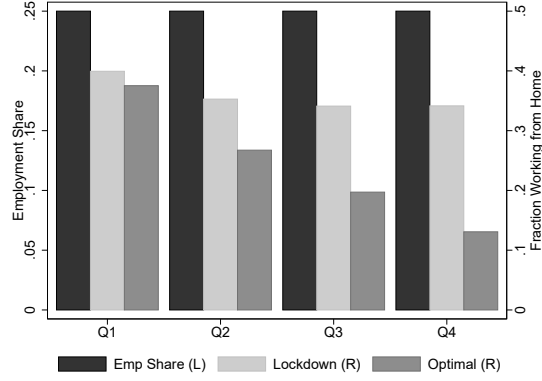
Distribution of Exposure Reduction and Wage Loss Because jobs are different in terms of exposure and WFH across occupations and industries, the reduction in exposure and the wage loss from the policies are distributed unequally across workers. In fact, low-wage jobs tend to have high exposure and low WFH, while high-wage jobs tend to have low exposure and high WFH. As a result, low-wage workers see large exposure reduction and wage losses under a lockdown, while the exposure reduction and wage losses are small for high-wage workers on average.¹³ Furthermore, since the real-world lockdown and the optimal policy target different jobs, the two policies have different distributional, as well as aggregate, consequences.

Figure 4 shows the distributional impacts of the actual lockdown and the optimal policy, across wage quartiles constructed from the average wage of occupation-industry combinations.¹⁴ In panel (a), the dark blue bars are the share of workers in each quartile, which is 25 percent by definition. The other two bars show, within each quartile, the fraction of workers working from home

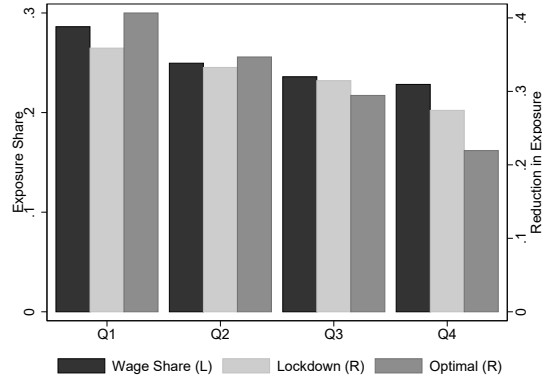
¹²Of course, there are other jobs that are truly essential. These can be incorporated as additional constraints to the minimization problem in (1).

¹³Aum et al. (2020a) document similar unequal impact of the pandemic even without a lockdown.

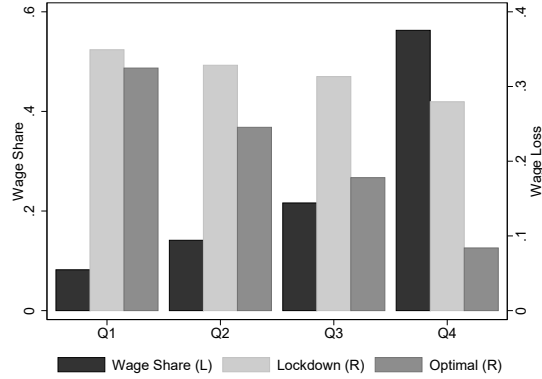
¹⁴The distributional impact of the constrained optimal policy is in the appendix.



(a) Fraction Working from Home



(b) Reduction in Exposure



(c) Wage Loss

Fig. 4: Policy Impact across Wage Quartiles

Each industry-occupation combination is ordered by its average wage and assigned to a quartile. In panel (a), the dark blue bars are each wage quartile's employment share, which is 0.25 by definition. The light blue and green bars depict the fraction of each quartile working from home due to the actual lockdown and the optimal policy, respectively, on the right scale. In panel (b), the dark blue bars are each wage quartile's share of aggregate exposure, and the light blue and the green bars depict the within-quartile reduction in exposure under the actual lockdown and the optimal policy, respectively. In panel (c), the dark blue bars are each wage quartile's share of aggregate wages, and the light blue and green bars depict within-quartile percentage wage losses by each policy.

due to the actual lockdown (light blue bars, right scale) and the optimal policy (green bars, right scale). The lockdown following the essential/non-essential distinction sends more low-wage workers home than high-wage workers (40 percent of the lowest wage quartile and 33 percent of the highest wage quartile). The difference is magnified under the optimal policy (38 percent of the lowest quartile and 12 percent of the highest quartile). There are two reasons. First, high-wage workers tend to have low exposure and are hence less likely to be sent home. Second, because of their high wage, holding other things equal, it is more costly to send home high-wage workers. Nevertheless, the optimal policy sends fewer workers home in all quartiles than the lockdown.

In panel (b), the dark blue bars depict each wage quartile’s share of the aggregate exposure, confirming that low-wage workers have higher exposure on average. The other two bars show the reduction in exposure due to the actual lockdown (light blue bars, right scale) and the optimal policy (green bars, right scale). Note that both policies reduce low-wage workers’ exposure more than that of the high-wage workers. The optimal policy strengthens this pattern, reducing low-wage workers’ exposure by more, but reducing high-wage workers’ exposure by less, compared to the lockdown.¹⁵ More significant, for the lower-wage quartiles, the optimal policy achieves a larger reduction in exposure while sending fewer workers home than under the actual lockdown.

In panel (c), the dark blue bars depict each wage quartile’s share of the aggregate wages, which is higher for higher quartiles by construction. The other two bars are the wage losses due to the lockdown (light blue bars, right scale) and the optimal policy (green bars, right scale). Both policies incur larger wage losses for lower-wage quartiles. This is because low-wage jobs are more likely to be sent home under both policies, but tend to be harder to perform from home (i.e., low WFH). The optimal policy leads to especially small wage losses for the top quartile, although it generates smaller wage losses than does the actual lockdown across all wage quartiles, as it sends home fewer workers across the board and also takes into account the wage losses.

In summary, low-wage workers are more affected than high-wage workers by both policies: They are more likely to be sent home, and as a result experience a larger reduction in exposure and larger wage losses. Comparing the two policies,

¹⁵This is by design, since we solved the optimal policy subject to the same reduction in aggregate exposure as the lockdown. The total reduction across all groups must be the same for the lockdown (light blue bars) and the optimal policy (green bars).

holding the reduction in aggregate exposure constant, low-wage workers see a larger reduction in exposure under the optimal policy, and high-wage workers a smaller reduction. However, although the optimal policy results in smaller wage losses for all wage quartiles compared to the lockdown, high-wage workers gain the most.

5 Conclusion

We construct indices of workers’ exposure to infection risks and work-from-home (WFH) that vary by both occupation and industry, and study their relationship across jobs. WFH varies widely even among jobs with similar levels of exposure, indicating that a planner could reduce the economic cost of a workplace lockdown by selectively sending home groups of workers based on the two indices, rather than using broad essential/non-essential categories. Compared to the actual lockdown, the optimal policy sends home one-third fewer workers and causes only half the loss in aggregate wages, while reducing aggregate exposure by the same magnitude. The wage gains are concentrated among high-wage workers, but low-wage workers benefit from a larger reduction in exposure, under the optimal policy.

While we abstract from some key dimensions, our work is a blueprint for an easily implementable smarter lockdown of the workplace during a pandemic. In addition, our cross-walk and merging of the O*NET and ATUS can be useful for a wider range of research, beyond the exposure and WFH indices used in this paper.

Bibliography

- Acemoglu, D. and D. Autor (2011). Chapter 12 - skills, tasks and technologies: Implications for employment and earnings. Volume 4, Part B of *Handbook of Labor Economics*, pp. 1043 – 1171. Elsevier.
- Adams-Prassl, A., T. Boneva, M. Golin, and C. Rauh (2020, Jun). Work that can be done from home: Evidence on variation within and across occupations and industries. IZA Discussion Paper 13374, Institute of Labor Economics (IZA), Bonn.
- Alon, T., M. Kim, D. Lagakos, and M. VanVuren (2020). How should policy responses to the Covid-19 pandemic differ in the developing world? Technical report, NBER.
- Assenza, T., F. Collard, M. Dupaigne, P. Fève, C. Hellwig, S. Kankanamge, and N. Werquin (2020, May). The Hammer and the Dance: Equilibrium and Optimal Policy during a Pandemic Crisis. CEPR Discussion Papers 14731, CEPR Discussion Papers.
- Aum, S., S. Y. T. Lee, and Y. Shin (2020a, May). Covid-19 doesn't need lockdowns to destroy jobs: The effect of local outbreaks in Korea. Working Paper 27264, National Bureau of Economic Research.
- Aum, S., S. Y. T. Lee, and Y. Shin (2020b, May). Inequality of fear and self-quarantines: Is there a trade-off between GDP and public health? Working Paper 27100, NBER.
- Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review* 103(5), 1553–97.
- Brotherhood, L., P. Kircher, C. Santos, and M. Tertilt (2020). An Economic Model of the Covid-19 Epidemic: The Importance of Testing and Age-Specific Policies. *Covid Economics - Vetted and Real-Time Papers*.
- del Rio-Chanona, R. M., P. Mealy, A. Pichler, F. Lafond, and J. D. Farmer (2020, 08). Supply and demand shocks in the COVID-19 pandemic: an industry and occupation perspective. *Oxford Review of Economic Policy* 36(Supplement_1), S94–S137.

- Dingel, J. I. and B. Neiman (2020). How many jobs can be done at home? *Journal of Public Economics* 189, 104235.
- Gamio, L. (2020, March). The workers who face the greatest coronavirus risk. The New York Times.
- Gottlieb, C., Grobovšek, M. Poschke, and F. Saltiel (2020). Lockdown accounting. *Covid Economics - Vetted and Real-Time Papers*.
- Hicks, M. J., D. Faulk, and S. Devaraj (2020, March). Occupational exposure to social distancing: A preliminary analysis using O*NET data. Technical report, Center for Business and Economic Research, Ball State University.
- Koren, M. and R. Pető (2020, 09). Business disruptions from social distancing. *PLOS ONE* 15(9), 1–14.
- Krueger, D., H. Uhlig, and T. Xie (2020, April). Macroeconomic dynamics and reallocation in an epidemic. *Covid Economics - Vetted and Real-Time Papers*.
- Leibovici, F., A. Santacreau, and M. Familgietti (2020, March). Social distancing and contact-intensive occupations.
- Mongey, S., L. Pilossoph, and A. Weinberg (2020, May). Which workers bear the burden of social distancing policies? *Covid Economics - Vetted and Real-Time Papers*.
- Palomino, J. C., J. G. Rodríguez, and R. Sebastian (2020). Wage inequality and poverty effects of lockdown and social distancing in Europe. *European Economic Review*, 103564.

Appendix

A Related Literature

A new strand of literature measures the degree to which jobs can be performed from home or are contact-intensive. One of the earlier papers is [Dingel and Neiman \(2020\)](#), which uses job characteristics in the O*NET to determine which occupations can be performed from home. [Koren and Pető \(2020\)](#), [Hicks et al. \(2020\)](#), and [Leibovici et al. \(2020\)](#) use O*NET to compute contact-intensity. Only few studies consider differences in such job characteristics across both industries and occupations. [Adams-Prassl et al. \(2020\)](#) collect information on WFH from geographically representative American and British surveys, and demonstrate a large WFH variation across industries even for the same occupation.

Most papers consider only either WFH or exposure separately. One exception is [Mongey et al. \(2020\)](#), which measures both physical proximity (one component of our exposure index) and WFH, and shows that there is a negative correlation between the two. Their WFH index is based on occupational characteristics from O*NET alone and does not vary industry, and they find a tighter correlation between physical proximity and WFH than what we find using the WFH based on ATUS, even when we ignore the industry dimension. [Adams-Prassl et al. \(2020\)](#) also document a negative relationship between their WFH index and O*NET physical proximity.

Few studies explicitly consider the costs of real-world lockdowns. [del Rio-Chanona et al. \(2020\)](#) construct an occupation-level remote labor index using O*NET and combine it with industry-wide lockdown measures to assess the heterogeneous effect of industry-level supply shocks across occupations. [Palomino et al. \(2020\)](#) construct a Lockdown Working Ability (LWA) index by combining a telework index from the O*NET and lockdown measures based on the government policies in Italy and Spain. They then simulate the impact of social distancing policies on inequality across European countries. [Gottlieb et al. \(2020\)](#) simulate the economic costs of various lockdown policies in developing countries, exploiting detailed data on each country’s demographic and labor market composition. [Aum et al. \(2020b\)](#) do the same for Korea and the United Kingdom but also model individuals’ choice of voluntarily working from home out of fear of infection.

To the best of our knowledge, ours is the only study that analyzes both WFH and exposure by industry and by occupation, and that introduces optimal

workplace social distancing policies based on these indices.¹⁶

B Alternative Policies

Exposure Reduction or Wage Loss The optimal policy minimizes aggregate wage losses subject to a given level of exposure reduction. We consider alternative programs that minimize exposure or wage losses subject to a given fraction of workers working from home.

The program that minimizes the aggregate exposure is:

$$\max_{\{x_i \in [0,1]\}_{i=1}^I} \psi \sum_{i=1}^I s_i e_i x_i \quad \text{s.t.} \quad \psi \sum_{i=1}^I s_i x_i \leq z, \quad (2)$$

which is in fact a maximization of the reduction in exposure subject to a given fraction of workers working from home. Clearly, the solution is to first send home those jobs with the highest exposure, regardless of the wage loss. The other program that minimizes aggregate wage losses subject to a given fraction of workers working from home is:

$$\min_{\{x_i \in [0,1]\}_{i=1}^I} \psi \sum_{i=1}^I s_i w_i (1 - h_i) x_i \quad \text{s.t.} \quad \psi \sum_{i=1}^I s_i x_i \geq z, \quad (3)$$

with the result that only jobs that earn a low wage and/or are easy to perform from home are sent home.

Figure 5 plots the reduction in exposure and aggregate wage loss from these two alternative policies for all levels of z (the fraction working from home) with $\psi = 0.487$. The outcomes of the optimal policy and the constrained optimal policy are also plotted for comparison. As shown in panel (a), unsurprisingly, the exposure minimization policy (line with diamonds) reduces aggregate exposure by more than any other policy, but at the cost of higher aggregate wage losses, as shown in panel (b). Conversely, the wage-loss minimization policy has the smallest aggregate wage loss of all policies, but at the cost of a smaller reduction in exposure than the other policies. The outcomes of the optimal policy lie between those of the two alternative policies.

¹⁶While most theoretical and structural papers that analyze the effect of the pandemic and lockdowns incorporate a trade-off between exposure and WFH (e.g. [Krueger et al., 2020](#); [Assenza et al., 2020](#)), few consider the heterogeneity of the trade-off at the micro level. Some exceptions are [Alon et al. \(2020\)](#) and [Brotherhood et al. \(2020\)](#), which consider differences by age, and [Aum et al. \(2020b\)](#), which considers differences by workers' occupation and skill level.

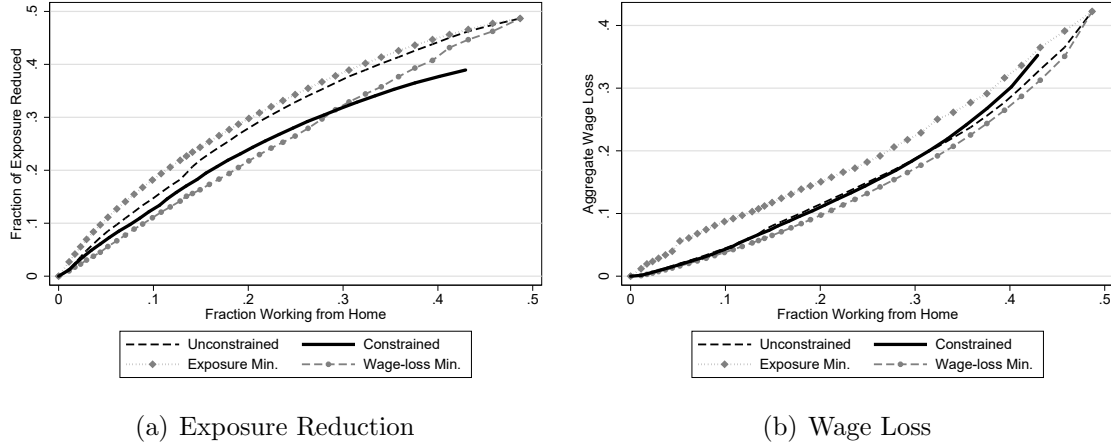


Fig. 5: Optimal Policy vs. Alternatives

In panel (a), for a given fraction of aggregate employment working from home on the horizontal axis, the reduction in aggregate exposure is plotted for the optimal policy (dashed line), the constrained optimal policy (solid line), the policy that minimizes aggregate exposure (line with diamonds), and the policy that minimizes aggregate wage losses (line with circles). Panel (b) plots the corresponding aggregate wage loss.

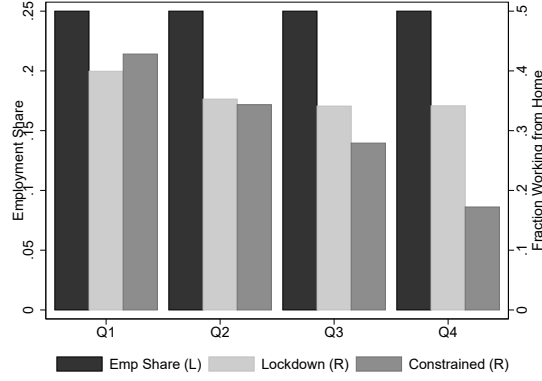
Impact of Constrained Optimal Policy across Wage Quartiles Finally, we show the counterpart of Figure 4 for the constrained optimal policy that keeps all health workers working normally.

The pattern across wage quartiles are similar to Figure 4, but there is one important difference. Relative to the actual lockdown, the constrained optimal policy sends home more bottom wage quartile workers, resulting in a larger wage loss for this group.

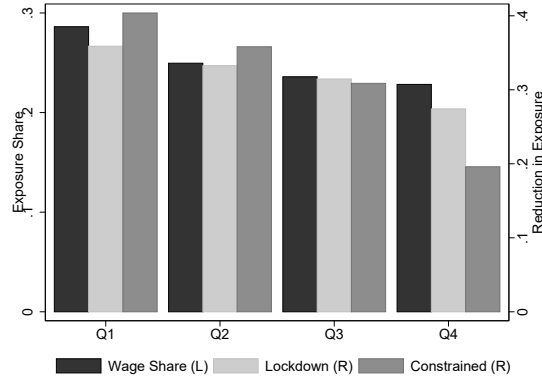
C Exposure and WFH by Demographic Group

How different demographic groups are affected by our index-based optimal policy depends on whether the variation in indices are captured by demographics. Thus, we regress the exposure index and the WFH index on demographic variables constructed for each occupation, industry, or industry-occupation pair: the female share, college share, Black and Hispanic shares, young and old shares, and the self-employment share.

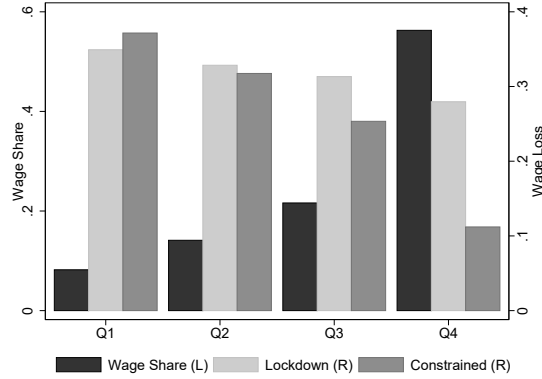
The results are shown in Table 2. Some group of workers, in particular college workers, are both less exposed to infection and less affected by working from home. But as expected from Figure 1, lower exposure does not necessarily mean a high WFH. Female, black and middle-aged workers tend to be more exposed to infection, but it is not related to how easily they can work from home. Hispanics



(a) Fraction Working from Home



(b) Reduction in Exposure



(c) Wage Loss

Fig. 6: Policy Impact across Wage Quartiles

Each industry-occupation combination is ordered by its average wage and assigned to a quartile. In panel (a), the dark blue bars are each wage quartile's employment share, which is 0.25 by definition. The light blue and green bars depict the fraction of each quartile working from home due to the actual lockdown and the constrained optimal policy, respectively, on the right scale. In panel (b), the dark blue bars are each wage quartile's share of aggregate exposure, and the light blue and the green bars depict the within-quartile reduction in exposure under the actual lockdown and the constrained optimal policy, respectively. In panel (c), the dark blue bars are each wage quartile's share of aggregate wages, and the light blue and green bars depict within-quartile percentage wage losses by each policy.

are less likely to work from home, but also face relatively lower risk of infection. In contrast, the self-employed work more from home, but do not necessarily face a lower risk of infection.

Demographics are correlated with both the exposure and WFH indices more by industry and less by occupation. And when combining industry and occupation together, demographics barely predict either index. This implies that industry-based lockdowns are at risk of disadvantaging vulnerable demographic groups, while more sophisticated policies that take into consideration both industry and occupation could reduce the likelihood of disproportionately affecting such groups.

Table 2: Regression Results

| | Occupation | | Industry | | OCC×IND | |
|----------|--------------------|--------------------|--------------------|-------------------|--------------------|--------------------|
| | exp | WFH | exp | WFH | exp | WFH |
| female | 1.24*** (0.43) | −0.02 (0.02) | 2.08*** (0.45) | −0.07* (0.04) | 1.11*** (0.22) | −0.00 (0.02) |
| ≥CLG | −0.88* (0.51) | 0.12*** (0.04) | −1.45*** (0.51) | 0.23*** (0.06) | −0.61* (0.33) | 0.13*** (0.03) |
| black | 4.32*** (1.12) | −0.03 (0.08) | 0.68 (0.76) | 0.11 (0.11) | 1.85*** (0.37) | −0.01 (0.04) |
| hispanic | −1.48* (0.81) | −0.23*** (0.08) | −1.50** (0.74) | −0.19* (0.11) | −0.67** (0.28) | −0.14*** (0.04) |
| young | −2.71*** (0.99) | 0.13 (0.08) | −4.71*** (1.08) | 0.22 (0.15) | −0.64 (0.41) | 0.05 (0.04) |
| old | −5.37*** (1.14) | 0.15* (0.09) | −5.15*** (1.29) | 0.13 (0.17) | −1.74*** (0.36) | 0.03 (0.03) |
| selfemp | 0.81 (0.49) | 0.29*** (0.07) | −0.64* (0.39) | 0.40*** (0.07) | 0.26 (0.24) | 0.26*** (0.05) |
| wage | 0.05 (0.57) | −0.00 (0.04) | −0.45 (0.42) | 0.02 (0.06) | 0.03 (0.29) | 0.01 (0.03) |
| const | 1.54 (2.12) | 0.01 (0.15) | 3.82** (1.63) | −0.12 (0.23) | 0.17 (0.99) | 0.02 (0.09) |
| R-sq | 0.320 | 0.448 | 0.606 | 0.510 | 0.222 | 0.159 |
| obs | 458 | | 254 | | 53,694 | |

Robust standard errors in parentheses. *, **, ***: significance at 90%, 95%, 99%.