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

Sangmin Aum, Sang Yoon (Tim) Lee, and Yongseok Shin

Shutting down the workplace is an effective means of reducing contagion but can induce large economic losses. We harmonize the American Time Use Survey and O*NET data to construct a measure of infection risk (exposure index) and a measure of the ease with which a job can be performed remotely (work-from-home index) across both industries and occupations. The two indexes are negatively correlated but distinct, so the economic costs of containing a pandemic can be minimized by sending home only those workers that are highly exposed to infection risk but that can perform their jobs easily from home. Compared with a lockdown of all non-essential jobs, which includes many jobs not easily performed from home, a more selective policy can attain the same reduction in aggregate infection risk (32 percent) with one-third fewer workers sent home to work (24 percent vs. 36 percent) and only half the aggregate wage loss (15 percent vs. 30 percent). In addition, moving to such a policy reduces the infection risk of low-wage workers the most and the wage losses of high-wage workers the most. Our crosswalk between the American Time Use Survey and O*NET data can be applied to a broader set of topics. (JEL E24, I14, J21)

1 INTRODUCTION

Most countries have implemented lockdowns and social distancing to varying degrees to contain the COVID-19 pandemic. The obvious downside is the economic costs, since most economic activities depend on in-person interaction. 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-to-infection risk from 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 article.

We begin by constructing an index of exposure-to-infection risk across occupations by using O*NET data and an index of how easily a job can be performed remotely (from home) across industries and occupations by using the "time worked from home" variable in the American Time Use Survey (ATUS). Such indexes are not new, but our approach is novel in that we harmonize the two datasets through our own crosswalk to quantify the trade-off between infection risk and the economic losses present in the economy, using the distribution of workers across jobs in the American Community Survey (ACS).

We find that, although jobs not easily performed from home (low work-from-home [WFH] ability) tend to be more exposed to infection risk on average, the negative correlation between the two indexes across occupations and industries is far from tight. First, there are a number of jobs with high exposure that can be easily performed from home (high WFH ability), such as IT sales agent. Second, infection risk varies widely even among jobs with the same WFH ability: For example, neither medical therapists nor experimental physicists can work from home, but the latter's work-places pose almost no risk of contagion. In addition, even the same occupation can have a very different WFH ability depending on the industry: For example, a registered nurse employed by a hospital has low WFH ability, but one in consulting services has high WFH ability.

In light of these findings, we consider an optimal policy that selectively sends home specific occupations in specific industries to minimize the aggregate wage loss subject to a given reduction in the aggregate exposure-to-infection risk. Intuitively, it is optimal to first send home workers with jobs with high exposure at work and small productivity and wage losses when working from home, the latter of which can be computed from a job's wage and WFH ability. Mathematically, this translates into a linear threshold in the two-dimensional plane of exposure and wage loss.

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 for half the economic cost. Under a constrained optimal policy in which healthcare-related workers must continue to work normally, the aggregate wage loss is 20 percent, still a third smaller than under a real-world lockdown. These gains are possible because the optimal policy exploits the large variation in WFH ability across occupations and industries for any given level of exposure—the novel fact we establish in this article.

It has become clear that low-wage workers are not only bearing the brunt of the pandemic economically but are also bearing the brunt of the infection risk. Compared with a lockdown of all non-essential jobs, the optimal policy reduces low-wage workers' infection risk the most. On the other hand, a move from such a 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 than when working from home. In fact, under all policy scenarios we consider, high-wage workers are the least exposed to infection risk and also lose the least economically, pointing to the importance of redistributive policies during a pandemic.

A word of caution: By design, our optimal policy is simple and static, abstracting from the essentiality of certain jobs (other than healthcare-related workers) that need to be performed even in the midst of a pandemic; the complementarity 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 higher WFH ability (to help prevent wage loss). It also assumes that the indexes we construct are constant, ignoring the potential change in exposure or WFH ability of specific jobs due to more subtle non-pharmaceutical interventions such as wearing masks, reorganizing the workplace, and changing individual behaviors as the population adjusts to the pandemic. However, it does capture the direct trade-offs given the work patterns at the onset of the pandemic. Furthermore, our analysis presents simple guidance for policymakers that is easy to implement in practice, while also providing a benchmark for structural economic models that consider some of the dimensions that we abstract from.[1]

It should also be noted that our industry-occupation crosswalk between ATUS and ACS/ O*NET is not specific to exposure and WFH ability and can be applied to a broader set of topics.

2 DATA

2.1 Employment and Wages by Occupation and Industry

We compute employment weights and mean hourly wages by occupation and industry from the ACS. We include only 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 and occupation codes from 2014 to 2018 as in Lee and Shin (2017). Finally, we take a simple average of the employment weights and hourly wages over the five years.

2.2 Constructing the Exposure and WFH Indexes

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 two questions: one about physical proximity (PP) to other people in the workplace and the other about potential exposure to disease or infection (EDI) in the workplace.[2]

We first convert O*NET titles to SOC codes using the accompanying crosswalk.[3] The SOC codes are then mapped to ACS OCC (occupation) codes using a crosswalk from 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 zero and standard deviation of 1, and take the average of the two as our exposure index value.[5]

WFH Index. Some earlier studies on COVID-19 used O*NET descriptors to construct a WFH index (e.g., Dingel and Neiman, 2020, and Mongey, Pilossoph, and Weinberg, 2020). However, a more accurate 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 the ATUS. While Mongey, Pilossoph, and Weinberg (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 indexes are misleading: For example, O*NET's "outdoor" categories (implying

Figure 1

Relationship Between Exposure and WFH Ability
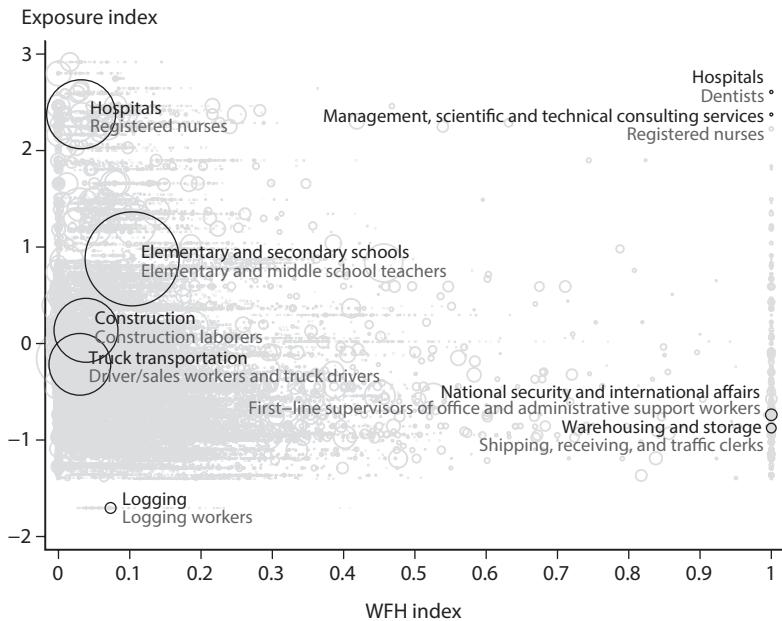
NOTE: Each circle represents a specific occupation within a specific industry (458 occupations × 254 industries). The size of the circle denotes the hours-adjusted employment share, averaged from 2014 to 2018. Black circles indicate examples of occupations either with very large employment shares or with extreme values of exposure and/or WFH ability. The first description for these examples is the industry, and the second is the occupation.

low WFH ability) include farmers who in fact have a high incidence of actually working from home, since they work in self-owned plots and land.[6] Quantitatively, the ATUS allows WFH ability to vary across industries, as well as across occupations, and indeed even the same occupation has very different WFH ability across industries in the data.[7]

For each industry-occupation combination in the ATUS, we compute the total time spent working and total time spent working 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 five years, and the ratio of the average time worked from home to the average time spent working is our WFH index value. Each cell is then matched to the ACS; in the process, we merge and impute missing cells using ACS employment weights and also ensure that each OCC code has at least one O*NET descriptor.

All in all, we end up with 458 occupations across 254 industries (458 × 254), each with an exposure index value, a WFH index value, hours worked, and hourly wages.

2.3 The Relationship Between Exposure and WFH Ability

Are jobs with higher infection risk 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 on the WFH index (horizontal axis) and exposure index (vertical axis). The size of the circle denotes the (hours-weighted) employment share of that job, averaged from 2014 to 2018.

Table 1

WFH Index

	Industry-occupation pairs	By occupation	By industry
Exposure	−0.033 (0.004)	−0.033 (0.005)	−0.045 (0.010)
Observations	54,108	458	254
R^2	0.034	0.109	0.126

NOTE: The table shows the result of regressing WFH ability on exposure, with observations weighted by their hours-adjusted employment shares. Robust standard errors are in parentheses. The units of observation are an industry-occupation pair (first column), occupation (second column), and industry (last column).

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

Table 1 shows the coefficients from regressing the WFH index on the exposure index for our whole industry-occupation sample (first column), only by occupation (second column), and only by industry (third column). For the latter cases, the exposure and WFH indexes are computed by averaging across each occupation's exposure and WFH ability within an industry, respectively, weighted by each occupation's within-industry employment share. In all three cases, the negative correlation between the exposure index and the WFH index 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 0.034 for the full sample of industry-occupation pairs. Since the exposure index varies only across occupations and not industries, by construction, the large increase in R^2 from the first to the second column mirrors the wide variation in WFH ability across industries—even for the same occupation.[8]

3 OPTIMAL POLICY

We interpret the fact that a job is more likely to be worked from home (high WFH ability) to mean that it will have little productivity loss when worked from home. Thus, the large dispersion in WFH ability among jobs with similar levels of exposure implies that the economic cost of sending home workers to reduce infection risk depends on whether their jobs have high or low WFH ability.

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

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

where ψ is a "compliance" constant that measures how many workers in the jobs sent home comply with the policy—that is, they work from home—and I is the total number of occupation-industry pairs.[9]

Designing the optimal policy is merely a high-dimensional linear programming problem.[10] The h_i in the objective function is our WFH index value that lies between 0 and 1. For the problem to be well defined, all e_i's need to be nonnegative, so we simply shift our exposure index by subtracting off its minimum value. 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.

The objective function is meant to capture the *aggregate* economic loss from sending home workers. As a result, the economic loss is larger not only when workers low on the WFH index are sent home, but also when industry-occupation pairs with a high wage and a large number of workers are sent home.[11] This finding does not necessarily mean that we are not interested in the distributive consequences of such policies. In fact, we discuss the unequal impact across workers in Section 4. We are merely separating the question of efficiency (i.e., minimizing the aggregate economic loss) from the question of redistribution, although we do not consider any redistributive policies in this article.

Figure 2 shows the optimal policy when $\psi = 0.473$.[12] In Panels A and B, each industry-occupation pair is plotted as a circle along the exposure (e_i, vertical axis) and the wage losses from sending home workers ($w_i(1 - h_i)$, 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 workers in jobs that are above the threshold are sent home (black circles). The slope of the threshold is positive since the optimal policy takes into account both exposure and wage losses from sending home workers. That is, even if a job has high infection risk, it is not sent home if it cannot be done remotely (low WFH ability) and has a high average wage, implying large wage losses from working remotely. Panel A is the solution when aggregate exposure is reduced by 10 percent and Panel B by 40 percent.

Panel C plots the fraction of workers sent home under the optimal policy (black lines with triangles) for each level of reduction in aggregate exposure (y) on the horizontal axis. Panel D shows the aggregate wage loss from the optimal policy (black lines with triangles against the same horizontal axis). Both black lines with triangles 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 the jobs with highest exposure and lowest wage losses first.

One unpleasant feature of the optimal policy we solve above is that many healthcare-related 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. For a more realistic problem during a pandemic, we also solve a constrained optimal policy in which healthcare-related workers must work normally. These are workers in the healthcare and social assistance sectors (industry codes 7970-8390) and healthcare practitioners, technical, and support occupations (OCC codes 3000-3655), who together comprise about 11 percent of (hours-adjusted) aggregate employment. Combined, these jobs account for about 20 percent of aggregate exposure.

Figure 2

Optimal Policy

A. Jobs working from home with $y = 0.1$

Exposure index

Wage loss

B. Jobs working from home with $y = 0.4$

Exposure index

Wage loss

C. Fraction working from home

Fraction working from home

Fraction of exposure reduced

D. Aggregate wage loss

Aggregate wage loss

Fraction of exposure reduced
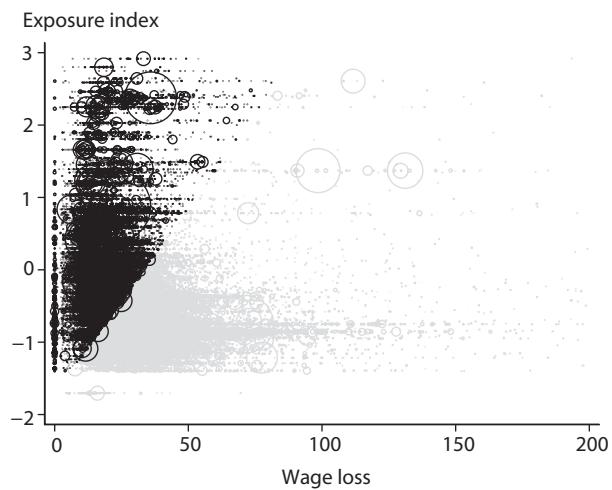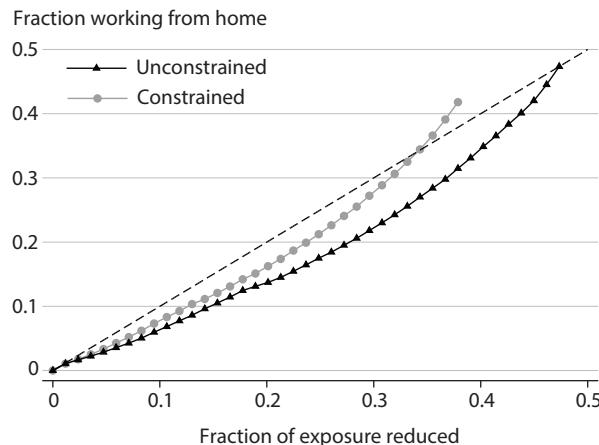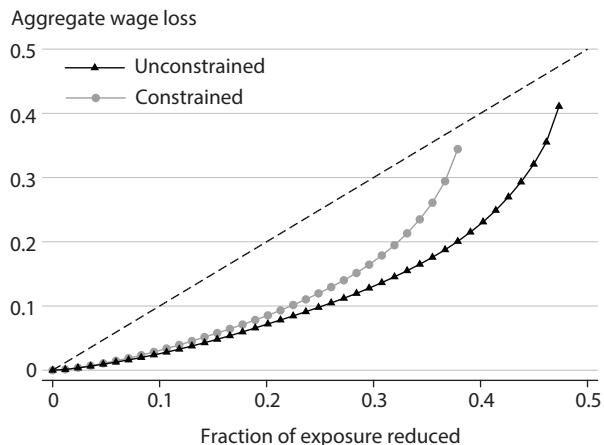
NOTE: In Panels A and B, each circle represents a specific occupation within a specific industry (458 occupations × 254 industries). The size of the circle denotes the hours-adjusted employment share, averaged from 2014 to 2018. The black circles represent jobs sent home by the optimal policy (group i such that $x_i^* = 1$), and the gray circles indicate jobs performed normally. Panel A is when the aggregate exposure is reduced by 10 percent and Panel B by 40 percent. In Panels C and D, for a fraction y of aggregate exposure reduction on the horizontal axis, the fraction of workers sent home and the aggregate wage loss (with the scaling constant $\psi = 0.473$ under the optimal policy) are shown as black lines with triangles in Panels C and D, respectively. The constrained optimal policy solves the same problem (1), but with the added constraint that all healthcare-related workers must work normally, as shown by the gray lines with circles.

In Panels C and D of Figure 2, the gray lines with circles plot the fractions of workers sent 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 with the unconstrained optimal policy, the constrained policy sends more workers home and causes higher wage losses for any given reduction in aggregate exposure, but it is still visibly convex.

The analysis is based on the pre-pandemic employment structure and wages of the United States as well as our exposure and WFH indexes for the United States. To the extent that the employment and wage distributions across occupations and industries differ across countries, the trade-off between infection and economic losses will differ across countries, even if countries share the same exposure and WFH indexes. In fact, trade-off patterns will also vary across regions or states within the United States.

4 OPTIMAL POLICY VS. REAL-WORLD LOCKDOWNS

We now compare our optimal policy with a lockdown that mimics those implemented in many U.S. states 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, governments classified industries as essential or non-essential and tried to keep essential workers working as normally as possible while also forbidding non-essential workers from commuting to work. For example, the Cybersecurity and Infrastructure Security Agency of the U.S. Department of Homeland Security provides guidelines on which jobs are essential for critical infrastructure.

Our version of the real-world lockdown follows Palomino, Rodríguez, and Sebastian (2020), who show which occupations and industries were effectively locked down in Europe.[13] In the context of our problem (1), jobs sent home in a real-world lockdown are a set $\{i \in \{1,2,\cdots,I\}|x_i = 1\}$, which in general will differ from the optimal solution. We set $\psi = 0.473$, as we did in the previous section, which generates a 30 percent drop in aggregate wages under the actual lockdown.[14] Given ψ, we find that a real-world lockdown reduces aggregate exposure by 32 percent.

Panel A of Figure 3 shows the non-essential jobs that sent workers home under the actual lockdown (dark-gray circles) and the jobs that sent workers home under the optimal policy (black circles). Both policies achieve the same reduction in aggregate exposure, but there is not much overlap between the two policies in terms of which jobs send workers home. In particular, the actual lockdown sent home many workers in jobs that have relatively low exposure. Panel B shows the jobs that sent workers home under the constrained optimal policy (black circles) (i.e., all healthcare-related workers working normally) still reduce aggregate exposure by 32 percent.

Table 2 shows exactly how this difference manifests in terms of the fractions of workers sent home and the aggregate wage loss. The gains from implementing the optimal policy are substantial: The same reduction in exposure can be attained by sending home one-third fewer workers (36 percent under the lockdown vs. 24 percent under the optimal policy) at half the economic cost (aggregate wage loss of 30 percent under the lockdown vs. 15 percent under the optimal policy).

Even with the constraint that all healthcare-related 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.[15]

Distribution of Exposure Reduction and Wage Loss. Because jobs are different in terms of exposure and WFH ability across occupations and industries, the reductions in exposure and in wage losses from the policies are distributed unequally across workers. In fact, low-wage jobs tend to have high exposure and low WFH ability, while high-wage jobs tend to have low exposure and high WFH ability. As a result, on average, low-wage workers see a large reduction in exposure and higher wage losses under a real-world lockdown, while high-wage workers see only a small reduction

Figure 3

Optimal Policy vs. Lockdown

A. Optimal

B. Constrained

● Lockdown   ● Unconstrained
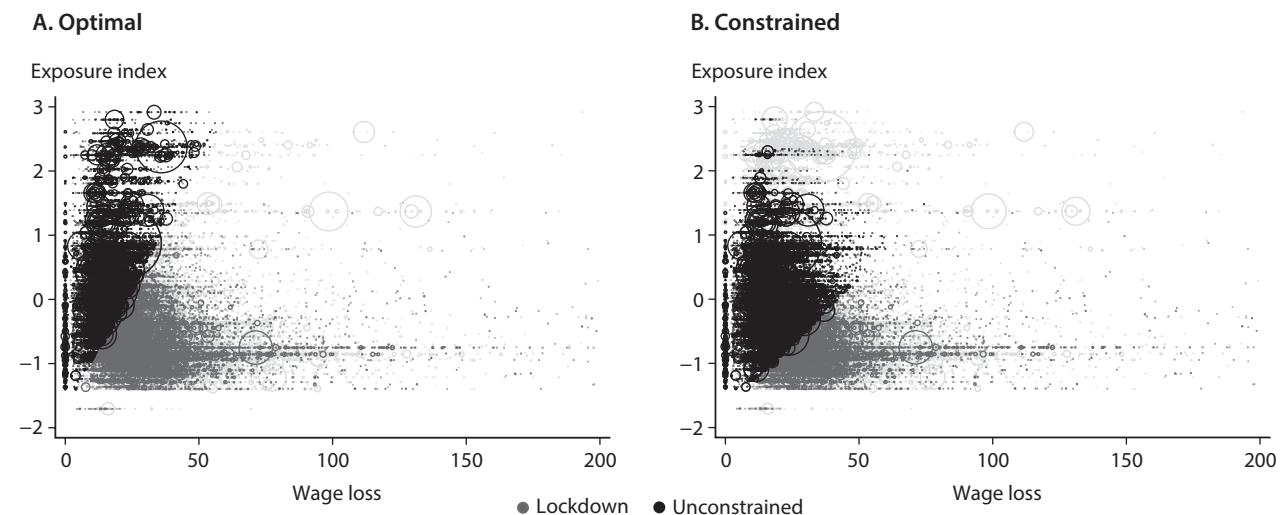
NOTE: Each circle is a specific occupation within a specific industry (458 occupations × 254 industries). The size of a circle represents its employment share. Dark-gray circles indicate jobs sent home during the actual lockdown. Black circles are the jobs sent home under the unconstrained optimal policy in Panel A and under the constrained optimal policy in the Panel B.

Table 2

Policy Outcomes

	Fraction sent home	Wage loss
A. Lockdown	0.359	0.300
B. Optimal	0.244	0.146
C. Constrained optimal	0.308	0.196

NOTE: The first column is the (hours-adjusted) fraction of workers sent home, and the second is the aggregate wage loss. In all three cases, aggregate exposure is reduced by 32 percent.

in exposure and lower wage losses.[16] 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 across the occupation-industry pairs.[17] In Panel A, the black bars are the shares of workers in each quartile, each 25 percent by definition. The other two bars within each quartile show the fractions of workers sent home under the actual lockdown (light-gray bars, right scale) and the optimal policy (dark-gray bars, right scale). The actual lockdown, which follows the essential/non-essential distinction, sends more low-wage workers home than high-wage workers (40 percent of the lowest wage quartile vs. 34 percent of the highest wage quartile). The difference is magnified under the optimal policy (37 percent of the lowest quartile vs.

Figure 4

Policy Impact Across Wage Quintiles

A. Fraction working from home
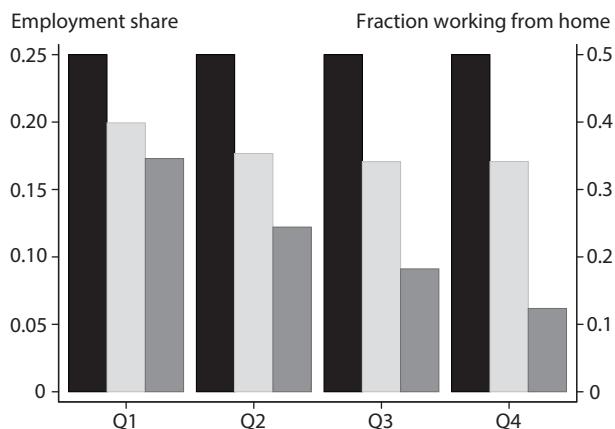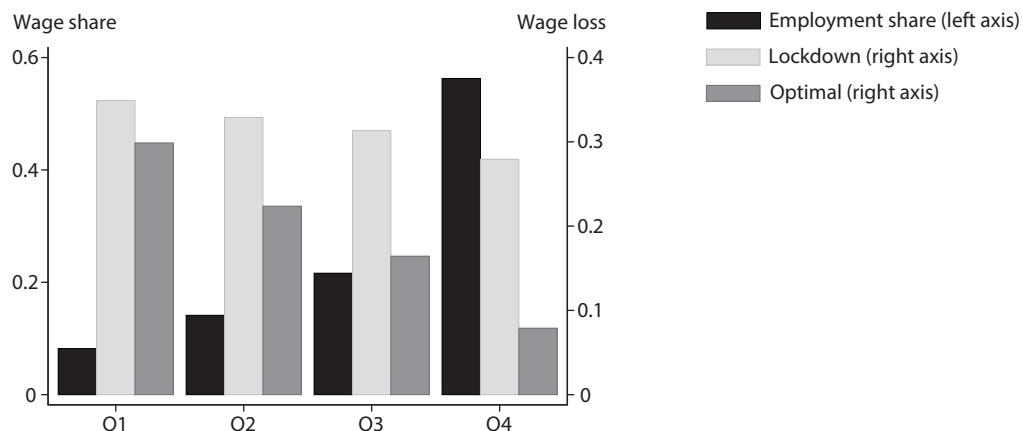
B. Reduction in exposure

C. Wage loss

NOTE: Each industry-occupation pair is ordered by its average wage and assigned to a quartile. In Panel A, the black bars are each wage quartile's employment share, which is 0.25 by definition. The light-gray and the dark-gray bars depict the fraction of each quartile sent home due to the actual lockdown and the optimal policy, respectively, on the right scale. In Panel B, the black bars are each wage quartile's share of aggregate exposure, and the light-gray and the dark-gray bars depict the within-quartile reduction in exposure under the actual lockdown and the optimal policy, respectively. In Panel C, the black bars are each wage quartile's share of the aggregate wage, and the light-gray and the dark-gray bars depict within-quartile wage losses under each policy.

13 percent of the highest quartile). There are two reasons for this difference. First, high-wage workers tend to have low exposure and are hence less likely to be sent home. Second, because of their high wages, 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 actual lockdown does.

In Panel B, the black bars are each wage quartile's share of 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-gray bars, right scale) and the optimal policy (dark-gray bars, right scale). Note that both policies reduce exposure more for low-wage workers than high-wage

workers. The optimal policy strengthens this pattern, reducing exposure more for low-wage than high-wage workers than the actual lockdown does.[18] More significant, for the lower wage quartiles, the optimal policy achieves a larger reduction in exposure while sending fewer workers home than the actual lockdown does.

In Panel C, the black bars are each wage quartile's share of aggregate labor income, which is higher for higher quartiles, by construction. The other two bars are the wage losses due to the actual lockdown (light-gray bars, right scale) and the optimal policy (dark-gray bars, right scale). Both policies incur larger wage losses for the lower wage quartiles. This is because low-wage jobs are more likely to send workers home under both policies but tend to be harder to perform from home (low WFH ability). The optimal policy leads to especially small wage losses for the top quartile, although it generates smaller wage losses than the actual lockdown does across all wage quartiles, as it sends home fewer workers across the board and also takes into account 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. As a result, they experience a larger reduction in exposure but also larger wage losses. Comparing the two policies, holding the reduction in aggregate exposure constant, low-wage workers see a larger reduction in exposure under the optimal policy, while high-wage workers see a smaller reduction. However, although the optimal policy results in smaller wage losses for all wage quartiles, high-wage workers lose the least relative to the actual lockdown.

5 CONCLUSION

We construct exposure and WFH indexes that vary by both occupation and industry and study their relationship across jobs. WFH ability 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 indexes, rather than using broad essential/non-essential categories. Compared with the actual lockdown, the optimal policy sends home one-third fewer workers and causes only half the losses in aggregate wages while also reducing aggregate exposure by the same magnitude. In addition, under the optimal policy, high-wage workers have the smallest wage losses but low-wage workers have the largest reduction in exposure.

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 crosswalk and merging of the ACS/O*NET and the ATUS can be applied to a wider range of research—beyond the exposure and WFH indexes used in this article. ∎

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 articles is by Dingel and Neiman (2020), who use job characteristics in the O*NET data to determine which occupations can work from home. Koren and Pető (2020), Hicks, Faulk, and Devaraj (2020), and Leibovici, Santacreau, and Familglietti (2020) use O*NET data to compute contact intensity. Only a few studies consider differences in such job characteristics across both industries and occupations. Adams-Prassl et al. (2020) collect information on WFH ability from geographically representative U.S. and British surveys and demonstrate large variation in WFH ability across industries even for the same occupation.

Most articles consider only either exposure or WFH ability. One exception is by Mongey, Pilossoph, and Weinberg (2020), who measure both PP (one component of our exposure index) and WFH ability and show 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 by industry, and they find a tighter correlation between PP and WFH ability than what we find using WFH ability based on the ATUS, even when we ignore the industry dimension. Adams-Prassl et al. (2020) also document a negative relationship between their WFH index and the O*NET PP data.

Few studies explicitly consider the costs of real-world lockdowns. Adams-Prassl et al. (2020) construct an occupation-level remote labor index using O*NET data and combine it with industry-wide lockdown measures to assess the heterogeneous effect of industry-level supply shocks across occupations. Palomino, Rodríguez, and Sebastian (2020) construct a lockdown working ability (LWA) index by combining a telework index from O*NET and lockdown measures based on 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, Lee, and Shin (2021a) do the same for Korea and the United Kingdom but also model the individual choice of voluntarily working from home out of fear of infection.

To the best of our knowledge, we are the first to analyze both WFH ability and exposure by industry and occupation and to go beyond merely documenting their correlation, by using indexes to determine which jobs to send home to minimize the aggregate wage loss for a given reduction in exposure.[19]

APPENDIX B: ALTERNATIVE POLICIES

B1 Exposure Reduction or Wage Loss

The optimal policy minimizes the aggregate wage loss subject to a given reduction in exposure. We consider alternative programs that minimize exposure or wage losses subject to a given fraction of workers sent home.

The program that minimizes the aggregate exposure is

(B1) $$\max_{\{x_i \in [0,1]\}_{i=1}^{I}} \psi \sum_{i=1}^{I} s_i e_i x_i \qquad \text{s.t.} \quad \psi \sum_{i=1}^{I} s_i x_i \leq z,$$

Figure B1

Optimal Policy vs. Alternatives

A. Exposure reduction

Fraction of exposure reduced

B. Wage loss

Aggregate wage loss

- - - - Unconstrained   ——— Constrained
········◆········ Exposure minimized   – · – · Wage-loss minimized
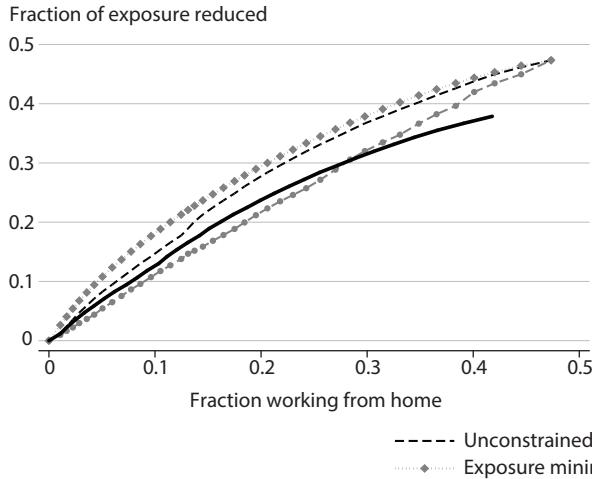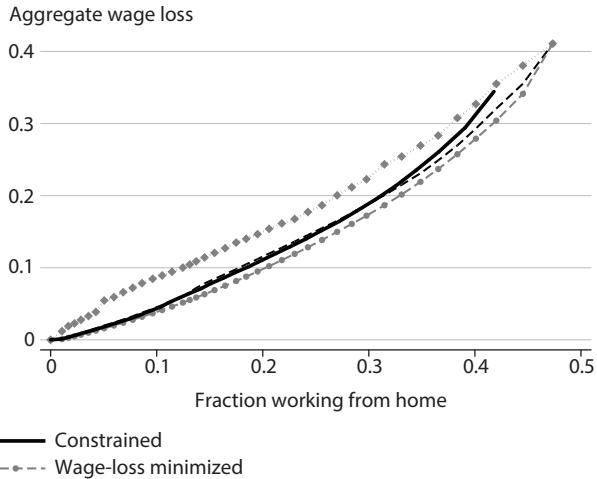
NOTE: In Panel A, for a given fraction of workers sent home on the horizontal axis, the reduction in aggregate exposure is plotted for the optimal policy (dashed line), the constrained optimal policy (solid line), the policy that minimizes aggregate exposure (line with diamonds), and the policy that minimizes the aggregate wage loss (line with circles). Panel B plots the corresponding aggregate wage loss.

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

$$\text{(B2)} \qquad \min_{\{x_i \in [0,1]\}_{i=1}^{I}} \psi \sum_{i=1}^{I} s_i w_i (1 - h_i) x_i \qquad \text{s.t.} \quad \psi \sum_{i=1}^{I} s_i x_i \geq z,$$

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

Figure B1 plots the reduction in exposure and the aggregate wage loss from these two alternative policies for all levels of z (the fraction of workers sent home) with $\psi = 0.473$. 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 loss, as shown in Panel B. Conversely, the wage-loss minimization policy (line with circles) 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.

B2 Impact of Constrained Optimal Policy Across Wage Quartiles

Finally, in Figure B2 we show the counterpart of Figure 4 for the constrained optimal policy that keeps all healthcare-related workers working normally.

Figure B2

Policy Impact Across Wage Quartiles

A. Fraction working from home
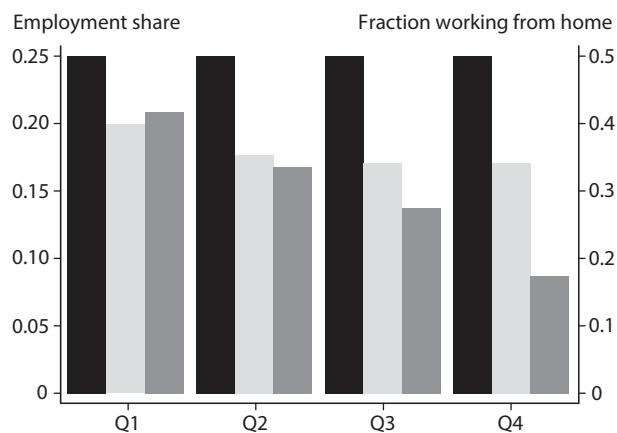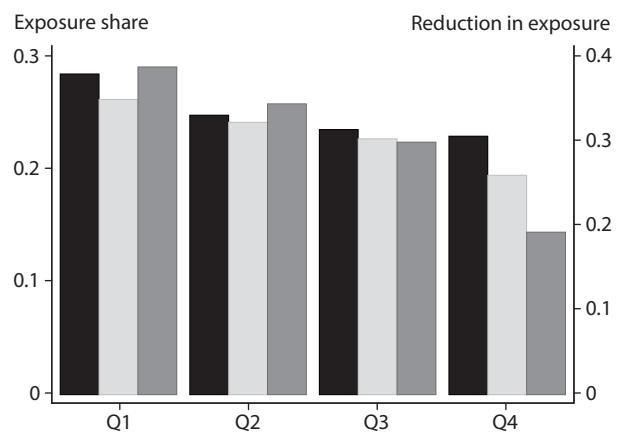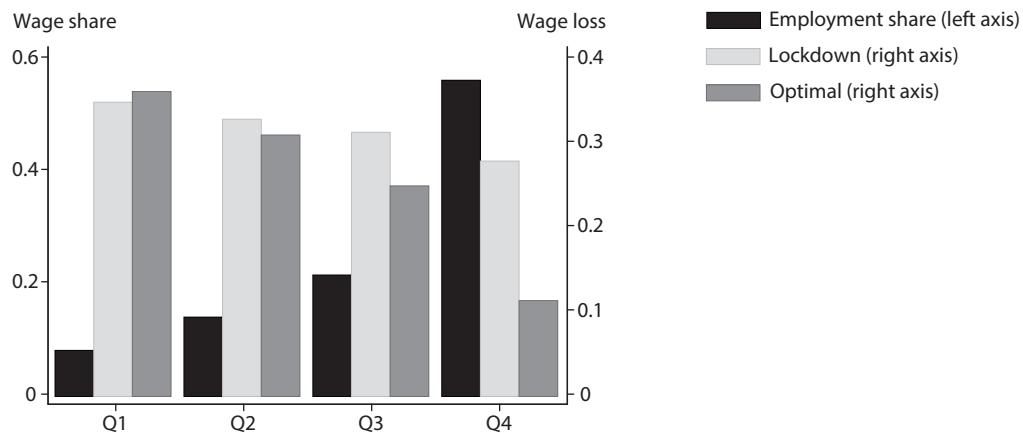
B. Reduction in exposure

C. Wage loss

■ Employment share (left axis)
▢ Lockdown (right axis)
■ Optimal (right axis)

NOTE: Each industry-occupation pair is ordered by its average wage and assigned to a quartile. In Panel A, the black bars are each wage quartile's employment share, which is 0.25 by definition. The light-gray and the dark-gray bars depict the fraction of each quartile sent home due to the actual lockdown and the constrained optimal policy, respectively, on the right scale. In Panel B, the black bars are each wage quartile's share of the aggregate exposure, and the light-gray and the dark-gray bars are the within-quartile reduction in exposure under the actual lockdown and the constrained optimal policy, respectively. In Panel C, the black bars are each wage quartile's share of the aggregate wage, and the light-gray and the dark-gray bars depict within-quartile wage losses by each policy.

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

Regression Results

	Occupation		Industry		Occupation-industry pair	
	Exposure	**WFH ability**	**Exposure**	**WFH ability**	**Exposure**	**WFH ability**
Females	1.24***	−0.02	2.08***	−0.07*	1.11***	−0.00
	(0.43)	(0.02)	(0.45)	(0.04)	(0.22)	(0.02)
≥ Bachelor's degree or more education	−0.88*	0.12***	−1.45***	0.23***	−0.61*	0.13***
	(0.51)	(0.04)	(0.51)	(0.06)	(0.33)	(0.03)
Blacks	4.32***	−0.03	0.68	0.11	1.85***	−0.01
	(1.12)	(0.08)	(0.76)	(0.11)	(0.37)	(0.04)
Hispanics	−1.48*	−0.23***	−1.50**	−0.19*	−0.67**	−0.14***
	(0.81)	(0.08)	(0.74)	(0.11)	(0.28)	(0.04)
Young (under 30 years of age)	−2.71***	0.13	−4.71***	0.22	−0.64	0.05
	(0.99)	(0.08)	(1.08)	(0.15)	(0.41)	(0.04)
Old (over 50 years of age)	−5.37***	0.15*	−5.15***	0.13	−1.74***	0.03
	(1.14)	(0.09)	(1.29)	(0.17)	(0.36)	(0.03)
Self-employed	0.81	0.29***	−0.64*	0.40***	0.26	0.26***
	(0.49)	(0.07)	(0.39)	(0.07)	(0.24)	(0.05)
Average hourly wage (in logs)	0.05	−0.00	−0.45	0.02	0.03	0.01
	(0.57)	(0.04)	(0.42)	(0.06)	(0.29)	(0.03)
Constant	1.54	0.01	3.82**	−0.12	0.17	0.02
	(2.12)	(0.15)	(1.63)	(0.23)	(0.99)	(0.09)
R^2	0.320	0.448	0.606	0.510	0.222	0.159
Observations	458		254		53,694	

NOTE: Robust standard errors are in parentheses. *, **, and *** indicate significance at the 90 percent, 95 percent, 99 percent levels, respectively.

APPENDIX C: EXPOSURE AND WFH ABILITY BY DEMOGRAPHIC GROUP

How different demographic groups are affected by our index-based optimal policy depends on whether the variations in the indexes 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 (under 30 years of age) and old (over 50 years of age) shares, and the self-employment share. The excluded groups in the regression are males, those without a four-year college degree, Whites, the middle aged (between 30 and 50 years of age), and employees.

The results are shown in Table C1. Some groups of workers, in particular college graduates, are both less exposed to infection risk and less likely to be sent home. But as expected from Figure 1, low exposure does not necessarily mean high WFH ability. Female, Black, and middle-aged workers tend to be more exposed to infection risk, but it is not related to how easily they can work from home. Hispanics are less likely to work from home but also face relatively lower infection risk. In contrast, the self-employed are more likely to work from home but do not necessarily face lower infection risk.

Demographics are correlated with both the exposure index and the WFH index more by industry and less by occupation. And when the industry-occupation pairs are used, demographics barely predict either index. This finding 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 those groups.

NOTES

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

[2] These 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 United States and reproduced for the United Kingdom by the U.K. Office for National Statistics. The *New York Times* and U.K. Office for National Statistics 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, Santacreau, and Famiglietti, 2020, and Mongey, Pilossoph, and Weinberg, 2020).

[3] Available at https://www.onetcenter.org/crosswalks.html.

[4] In general, the O*NET SOC codes are finer than the ACS OCC codes, 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, O*NET does not list descriptors for any lower level of some of the three-digit OCC codes, and OCC codes changed in 2018, necessitating additional manipulation by us.

[5] PP and EDI are positively correlated. In particular, the highest EDI occupations (mostly healthcare-related occupations) also have high PP. At the same time, there are occupations that have high EDI but low PP (e.g., couriers, refuse collectors, and janitors and building cleaners) and those that have high PP but low EDI (e.g., meeting and event planners, engine and other machine assemblers, and parts salespersons).

[6] Bick, Blandin, and Mertens (2020) also emphasize the difference between index-based potential home-based work and actual home-based work.

[7] Adams-Prassl et al. (2020) show that this is also true in their own surveys in the United States and the United Kingdom.

[8] Appendix C shows that the correlation between our indexes and demographics are also weak.

[9] Thus, the maximal possible reduction in aggregate exposure is ψ, which may be due to non-compliance, selective furloughs, or a 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 the magnitudes of the reductions in wage loss and exposure.

[10] Problem (1) is the dual problem of minimizing exposure subject to a given level of the aggregate wage loss.

[11] We are not considering job losses at the extensive margin. One could assume that a WFH index h_i below a certain threshold $\bar{h}$ implies workers sent home are too unproductive to be employed and hence lose their jobs. With this assumption, one could work out an alternative problem that minimizes job losses rather than the aggregate wage loss.

[12] $\psi = 0.473$ corresponds to a 30 percent drop in aggregate wages under the real-world lockdown, to be shown in Section 4.

[13] Palomino , Rodríguez, and Sebastian (2020) identify essential jobs by ISCO (two digit) and NACE (one 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 (Cybersecurity and Infrastructure Security Agency) guidelines.

[14] A labor income share of 60 percent implies an 18 percent drop in gross domestic product (GDP) due to the lockdown, which is between the GDP loss in the United States (10 percent) and the GDP loss in the United Kingdom (20 percent) from the first to the second quarter of 2020.

[15] Of course, there are other jobs that are truly essential. Those jobs can be incorporated as additional constraints to the minimization problem in (1). While the exemption of a larger set of truly essential jobs will bring our constrained policy closer to real-world lockdowns, one important distinction is that our constrained policy is at the level of industry-occupation pairs, while the real-world policies were industry based and hence more blunt.

[16] Aum, Lee, and Shin (2021b) document a similar unequal impact of the pandemic—even without a lockdown.

[17] The distributional impact of the constrained optimal policy is in the appendix.

[18] This is by design, since we solved the optimal policy subject to the same reduction in aggregate exposure as for the lockdown. The total reduction across all groups must be the same for the lockdown (light-gray bars) and the optimal policy (dark-gray bars).

[19] While most theoretical and structural articles that analyze the effect of the pandemic and lockdowns incorporate a trade-off between exposure and WFH ability (e.g., Krueger, Uhlig, and Xie, 2020, and 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), who consider differences by age, and Aum, Lee, and Shin (2021a), who consider differences by worker occupation and skill level.

REFERENCES

Adams-Prassl, A; Boneva, T.; Golin, M. and Rauh, C. "Work That Can Be Done from Home: Evidence On Variation Within and Across Occupations and Industries." IZA Discussion Paper 13374, Institute of Labor Economics, June 2020.

Alon, T.; Kim, M.; Lagakos, D. and VanVuren, M. "How Should Policy Responses to the COVID-19 Pandemic Differ in the Developing World?" NBER Working Paper 27273, National Bureau of Economic Research, May 2020; https://doi.org/10.3386/w27273.

Assenza, T.; Collard, F.; Dupaigne, M.; P. Féve; Hellwig, C.; Kankanamge, S. and Werquin, N. "The Hammer and the Dance: Equilibrium and Optimal Policy during a Pandemic Crisis." CEPR Discussion Paper 14731, Centre for Economic Policy Research, May 2020; https://cepr.org/active/publications/discussion_papers/dp.php?dpno=14731.

Aum, Sangmin; Lee, Sang Yoon (Tim) and Shin, Yongseok. "Inequality of Fear and Self-Quarantine: Is There a Trade-Off Between GDP and Public Health?" *Journal of Public Economics*, February 2021a, Volume 194; https://doi.org/10.1016/j.jpubeco.2020.104354.

Aum, Sangmin; Lee, Sang Yoon (Tim) and Shin, Yongseok. "COVID-19 Doesn't Need Lockdowns to Destroy Jobs: The Effect of Local Outbreaks in Korea." *Labour Economics*, June 2021b, Volume 70; https://doi.org/10.1016/j.labeco.2021.101993.

Autor, D.H. and Dorn, D. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, August 2020, *103*(5), 2013, pp. 1553-97; https://doi.org/10.1257/aer.103.5.1553.

Bick, A.; Blandin, A. and Mertens, K. "Work from Home after the COVID-19 Outbreak." Unpublished manuscript, 2020.

Brotherhood, Luiz; Kircher, Philipp; Santos, Cezar and Tertilt, Michèle. "An Economic Model of the Covid-19 Epidemic: The Importance of Testing and Age-Specific Policies." CESifo Working Paper Series 8316, 2020; http://dx.doi.org/10.2139/ssrn.3618840.

del Rio-Chanona, R.M.; Mealy, P.; Pichler, A.; Lafond, F. and Farmer, J.D. "Supply and Demand Shocks in the COVID-19 Pandemic: An Industry and Occupation Perspective." *Oxford Review of Economic Policy*, August 2020, *36*(Supplement 1), pp. S94-S137; https://doi.org/10.1093/oxrep/graa033.

Dingel, J.I. and Neiman, B. "How Many Jobs Can Be Done at Home?" *Journal of Public Economics*, September 2020, *189*, 104235; https://doi.org/10.1016/j.jpubeco.2020.104235.

Gamio, L. "The Workers Who Face the Greatest Coronavirus Risk." *New York Times*, March 15, 2020.

Gottlieb, C.; Grobovšek, J; Poschke, M. and Saltiel, F. "Lockdown Accounting." Centre for Economic Policy Research *Covid Economics: Vetted and Real-Time Papers*, June 2020.

Hicks, M.J.; Faulk, D. and Devaraj, S. "Occupational Exposure to Social Distancing: A Preliminary Analysis Using O*NET Data." Technical report, Center for Business and Economic Research, Ball State University, March 2020.

Koren, M. and Pető, R. "Business Disruptions from Social Distancing." *PLOS ONE*, September 18, 2020, *15*(9), pp. 1-14; https://doi.org/10.1371/journal.pone.0239113.

Krueger, D.; Uhlig, H. and Xie, T. "Macroeconomic Dynamics and Reallocation in an Epidemic." Centre for Economic Policy Research *Covid Economics: Vetted and Real-Time Papers*, April 2020; https://cepr.org/sites/default/files/CovidEconomics5.pdf.

Lee, S.Y. and Shin, Y. "Horizontal and Vertical Polarization: Task-Specific Technological Change in a Multi-Sector Economy." NBER Working Paper 23283, National Bureau of Economic Research, March 2017; https://doi.org/10.3386/w23283.

Leibovici, F.; Santacreau, A. and Familglietti, M. "Social Distancing and Contact-Intensive Occupations." Federal Reserve Bank of St. Louis *On the Economy Blog*, March 24, 2020; https://www.stlouisfed.org/on-the-economy/2020/march/social-distancing-contact-intensive-occupations.

Mongey, S.; Pilossoph, L. and Weinberg, A. "Which Workers Bear the Burden of Social Distancing Policies?" Centre for Economic Policy Research *Covid Economics: Vetted and Real-Time Papers*, May 2020; https://cepr.org/sites/default/files/CovidEconomics12.pdf.

Palomino, J.C.; Rodríguez, J.G. and Sebastian, R. "Wage Inequality and Poverty Effects of Lockdown and Social Distancing in Europe." *European Economic Review*, 2020, *129*, 103564; https://doi.org/10.1016/j.euroecorev.2020.103564.