NBER WORKING PAPER SERIES

# A SOCIAL INSURANCE PERSPECTIVE ON PANDEMIC FISCAL POLICY: IMPLICATIONS FOR UNEMPLOYMENT INSURANCE AND HAZARD PAY

Christina D. Romer
David H. Romer

Working Paper 29419
http://www.nber.org/papers/w29419

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2021

This paper is in preparation for the *Journal of Economic Perspectives*. We are grateful to Peter Ganong, Eric Hurst, Nina Pavcnik, Timothy Taylor, and Heidi Williams for helpful comments, and to Jacob Weber for research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

A Social Insurance Perspective on Pandemic Fiscal Policy: Implications for Unemployment Insurance and Hazard Pay
Christina D. Romer and David H. Romer
NBER Working Paper No. 29419
October 2021
JEL No. E62,E65,H21,H31,I18

## ABSTRACT

This paper considers fiscal policy during the pandemic through the lens of optimal social insurance. We develop a simple framework to analyze how government taxes and transfers could mimic the insurance against pandemic income losses that people would like to have had. Permutations of the framework provide insight into how unemployment insurance should be structured, when and how much hazard pay is called for, and whether fiscal policy should aim just to redistribute income or also to stimulate aggregate demand during a pandemic. When we use the insights from the model to evaluate unemployment insurance measures taken during the pandemic, we find that some, but far from all, of the implications of the social insurance framework were followed. In the case of hazard pay, we find that the proposal for a national program (the never-implemented HEROES Act) was both broader and more generous than a social insurance perspective would call for. We suggest that the social insurance perspective on fiscal policy is likely to become increasingly relevant as pandemics and climate-related natural disasters become more common causes of unemployment and recessions.

Christina D. Romer
Department of Economics
University of California, Berkeley
Berkeley, CA 94720-3880
and NBER
cromer@berkeley.edu

David H. Romer
Department of Economics
University of California, Berkeley
Berkeley, CA 94720-3880
and NBER
dromer@berkeley.edu

Covid-19 is a worldwide public health crisis unlike anything seen in the post-World War II era. The fiscal response has been similarly extraordinary. The United States has spent $5.2 trillion on a wide range of recovery and support initiatives, and other countries have also spent unprecedented sums. The pros and cons of various measures have been discussed extensively. But what has been largely missing from the fiscal policy discussion is an overarching analytical framework that takes into account the unique nature of a pandemic recession. This paper seeks to fill this gap.

We argue that a social insurance perspective is the appropriate way of understanding and evaluating fiscal policy in a pandemic. Social insurance analysis shows how the government can use taxes and transfers to provide people with insurance that they would like to have, but that doesn't exist or that no one had contemplated needing. During a pandemic, workers in certain sectors face prolonged unemployment because their industries can't operate safely, while workers in other sectors remain relatively unscathed. Had workers foreseen this possibility, they would have liked to purchase insurance against the risk that their sector would be closed. The social insurance framework can show which types of government fiscal actions best approximate what a well-functioning insurance market would provide.

A social insurance perspective is more appropriate for designing and evaluating pandemic fiscal policy than simple aggregate-demand-based models. Conventional Keynesian models of fiscal policy suggest that the way to deal with a recession is to increase aggregate demand quickly, and by enough to return output to its normal or potential level. And in this framework, it is not necessary for fiscal policy to closely target the workers or industries most affected by the recession. Raising aggregate demand anywhere will raise incomes and spending throughout the economy, and so help will eventually flow to those most affected.

These models and policy prescriptions don't hold in a pandemic recession. Because the virus thrives on human interaction (and hence on some types of economic activity), fiscal policy should not be aimed at quickly raising aggregate demand and attempting to return the economy to full employment. Doing so would make the pandemic worse and increase illness and deaths. Similarly, in a pandemic, some types of economic activity—such as in-restaurant dining and cruise travel—simply can't take place safely. As a result, broad stimulus measures like one-time payments or tax cuts can do little to put workers in those industries back to work. In essence, pandemic-related shutdowns of certain sectors short-circuit the usual Keynesian multiplier effect.

We begin our analysis by developing the social insurance framework. We describe a simple model with one sector that's affected by the pandemic and one that isn't. We show that optimal policy involves the government taxing those in the unaffected sector and providing income support for those in the sector that's shut. We then discuss several enhancements of the framework that yield richer understanding and more nuanced policy prescriptions. For example, adding a third sector where essential workers remain employed, but face greater health risks because of the nature of their jobs, suggests that government-provided hazard pay is appropriate. Or, incorporating notions of fairness or difficulty in identifying which workers remain able to work results in optimal policy that includes a role for general stimulus. The appendix to the paper contains a more formal presentation of the model.

Armed with the framework, we examine two types of pandemic fiscal policy in detail, one that's been used extensively—unemployment insurance—and one that's hardly been used at all—hazard pay. We discuss in more depth the implications of a social insurance perspective for the usefulness of such policies and how they should be structured. We also discuss some of the practical issues around designing these fiscal actions for use in a pandemic, and examine how the implications of a social insurance perspective compare with what was actually enacted or proposed during the pandemic.

Finally, in the conclusion, we consider the broader applicability of the social insurance perspective on fiscal policy. We argue that the insights likely carry over to a wide range of situations other than the current pandemic.

Some previous authors have also suggested a social insurance perspective on the appropriate policy response to the pandemic, and the baseline case of the next section draws heavily on that work.[1] Our contributions are in the ways we go beyond the baseline case and in evaluating actual and proposed policies from a social insurance perspective.

## A Social Insurance Perspective

The basic logic of a social insurance perspective on the fiscal policy response to a pandemic is straightforward. We therefore present a verbal description here, and leave the mathematical analysis to the appendix.

### A Baseline Case

To see many of the key implications of the social insurance perspective, we start with a very stylized case that builds on Guerrieri et al. (2020) and Woodford (2020). We then consider the implications of relaxing some of its assumptions to address additional issues.

Think of a competitive, one-period economy with many identical individuals. There are two sectors, *A* and *B*. Other than producing different outputs, the two sectors are identical. Each individual obtains utility from their consumption of both sectors' outputs. Consumption of one sector's output doesn't affect utility from consuming the other sector's (that is, utility is "additively separable" over consumption of the two outputs). Although this assumption of course isn't exactly correct, the idea that, for example, whether one is able to go the dentist or take an airplane flight doesn't affect the marginal utility of groceries or clothes is arguably a reasonable approximation.

[1] In a contribution early in the pandemic, Milne (2020) argues that the idea of "retrospective insurance" provides a valuable way of thinking about important aspects of the appropriate policy response. Similarly, in an informal early contribution, Saez and Zucman (2020) propose full government replacement of lost income for workers and businesses as a form of social insurance. Our framework builds most closely on Woodford (2020) and, especially, Guerrieri et al. (2020). Both papers consider multi-sector models where one sector is forced to shut because of a pandemic, and both consider social insurance policies. And like us, Woodford uses the hypothetical case where there are perfectly functioning markets for "pandemic insurance" as a benchmark. These authors' main interests, however, are different from ours: Guerrieri et al.'s is in the conditions under which a pandemic—which is fundamentally a shock to aggregate supply—can lead to an aggregate demand shortfall (an issue we discuss below), and Woodford's is in the consequences of the structure of linkages among sectors.

We also assume that individuals get disutility if they work. For simplicity, work is a 0-1 variable—an individual is either working or they aren't.

We think of a pandemic as an event that makes production in one sector impossible; this could be either because producing the sector's output is unsafe or because consuming it is. To preserve the symmetry between the sectors, we assume they have the same probability of being shut by a pandemic.

Finally, individuals must decide which sector they want to work in before they learn whether a pandemic will occur. Allowing for some ex post mobility mitigates the fall in output in a pandemic, but otherwise has little impact on the messages of the baseline case.

The equilibrium of this stylized economy isn't hard to describe. Consider first what happens when individuals cannot insure against a pandemic. This could occur because people simply hadn't contemplated the possibility of a pandemic, or because in practice the difficulty of spelling out exactly what constitutes a pandemic makes the cost of writing insurance contracts prohibitive. The symmetry of the model implies that half of individuals are in each sector. In the absence of a pandemic, everyone earns the same amount, each sector produces the same amount, and each individual consumes the same amount of each sector's output. If a pandemic shuts one sector, the individuals who were working there earn no income, and so have no consumption (recall that the economy lasts for only one period and workers can't switch sectors in a pandemic). The individuals in the sector that stays open continue to work, but they now spend all their income on that sector's output rather than splitting it between the two sectors.

In normal times, because everyone's consumption is the same, everyone's marginal utility of consumption is the same. But in a pandemic, the marginal utility of consumption (of the output of the sector that remains open) of those who remain employed is lower than normal, while the marginal utility of those who become unemployed is higher than normal (probably greatly so, since their consumption falls to zero). Thus, the marginal utility of sector-*A* workers is higher than that of sector-*B* workers if there's a sector-*A* pandemic, but lower if there's a sector-*B* pandemic.

This variation in relative marginal utilities implies that from an ex ante perspective—that is, before it's known whether there will be a pandemic—the outcome is Pareto inefficient. Measures that would shift resources from low-marginal-utility individuals to high-marginal-utility individuals in a pandemic would raise everyone's ex ante expected utility.

The efficient allocation can be achieved without any government action if there are not only competitive markets for output and labor in each sector after it's known whether there will be a pandemic, but also competitive markets for "pandemic insurance" before it's known whether there will be a pandemic. With these markets in place, the outcome in the absence of a pandemic is the same as before. But now individuals in each sector purchase insurance against the possibility of a pandemic hitting their sector (with the market clearing because the individuals in the other sector are willing to sell such insurance). In the event of a pandemic, every individual consumes the same amount of the output of the sector that stays open as they do in the absence of a pandemic. Marginal utility is always equal across all individuals, and so the allocation is efficient.[2]

The final step is to return to the case where there's no insurance and consider how the government can implement the efficient allocation—that is, the one that would occur with perfectly functioning insurance markets. In this case, the public sector is providing the insurance individuals would obtain for themselves if markets functioned perfectly, so it's natural to describe the policy as social insurance.

The policy that reproduces the efficient outcome is simple. In the absence of a pandemic, the government takes no action. But if there is a pandemic, it taxes the individuals in the sector that is still open half their income and transfers the proceeds to the individuals in the sector that is

[2] Another market-based way of achieving the efficient allocation is through trade in competitive markets in the full set of "state-contingent" (or "Arrow-Debreu") commodities (output and labor in each sector in each state of the world—no pandemic, a sector-*A* pandemic, and a sector-*B* pandemic) before it's known whether there will be a pandemic. This produces the same outcomes as with competitive markets for pandemic insurance. Although this set-up is obviously more abstract and even less realistic than the possibility of pandemic insurance, it can be useful for clarifying outcomes and their welfare properties.

shut. The result is that each individual's after-tax-and-transfer income is half their usual income, and everyone's consumption of the output of the sector that stays open is the same as in normal times.

Thus, the policy that replicates what would happen with perfect insurance markets is one of targeted transfers: the government makes transfers to individuals who can't work because of the pandemic. Moreover, there's no "stimulus" in the optimal policy: what the government pays out in transfers equals what it takes in from taxes, and output and employment are the same as they would be without insurance markets or government intervention (as also observed by Woodford, 2020). In addition, although individuals who can't work because of the pandemic obtain the same after-tax-and-transfer income (and the same consumption) as those who stay employed, they don't attain their usual income. Instead, they only get what they normally spend on the output of the non-pandemic sector. In that sense, insurance is less than complete.

Finally, a simple way of thinking about the government's policy is that it taxes the individuals in the non-pandemic sector what they would normally spend on the pandemic sector's output, and gives the proceeds to the individuals in the pandemic sector. This allows them to maintain their normal spending on the output of the non-pandemic sector.

Although this baseline case shows some key messages of the social insurance perspective, it omits some important issues. We therefore turn to extensions. Table 1 summarizes their key features and implications.

## Incentives and Fairness

In the efficient allocation of the baseline case, the individuals who were working in the sector that shuts down are actually better off than those in the sector that stays open: everyone has the same consumption, but only the individuals in the sector that's open have the disutility of working. There are two reasons that implementing that allocation through government social insurance might not be workable. The first involves incentives. In the model, it's evident who should continue working in the pandemic—everyone working in the sector that isn't shut by the pandemic. In practice, however, individuals differ in their nonwork situations (such as their childcare needs and preexisting health conditions) and in their attitudes toward health risks; some individuals can change sectors fairly easily in a pandemic; and it often isn't obvious what firms (or what parts of a firm) can continue operating in a pandemic. As a result, for many workers it's not possible to tell whether the optimal outcome involves them continuing to work. Second, public perceptions of fairness don't align perfectly with allocations that would prevail under complete markets (see, for example, Weinzierl, 2014), so an allocation that makes the unemployed better off than the employed may not be politically feasible or may be viewed as undesirable.

We capture these considerations by assuming the government has no information about who should optimally be working in a pandemic and who shouldn't. As a result, the individuals who continue working must be at least as well off as those who don't; otherwise, no one would work. Since working involves disutility, it follows that the employed must have greater consumption than the unemployed. Total output (and hence total consumption) is the same as without this constraint. Thus the employed consume more than they do without the constraint, and the unemployed consume less. The "replacement rate" from social insurance is therefore lower than in the baseline case. The other key implications of the baseline case continue to hold: the optimal policy is targeted transfers, and there's no need for stimulus.[3]

### High-Risk Essential Workers and Hazard Pay

During a pandemic, there are some jobs—notably in healthcare—where it's socially desirable for workers to continue working even though it involves elevated risk of illness. To capture this,

[3] The appendix also discusses the cases where the government has some but less than full information about who should be working in a pandemic, and where some individuals can switch sectors in a pandemic but the government has limited information about who can switch easily. For the most part, these extensions don't change our main messages. But we find that labor mobility with limited government information not only leads to a smaller fall in output in a pandemic than in the baseline case, but also provides a reason the government might want to provide a "moving bonus" to workers who switch sectors in a pandemic (somewhat analogous to the possibility of hazard pay discussed below).

suppose there are three sectors rather than two. When there's a pandemic, as before one sector shuts down and one is unaffected. But now a third sector is essential and high-risk: the workers are particularly valuable during a pandemic but must interact with others, and so face greater chances of illness. As a result, both the utility of consuming the sector's output and the disutility of working in this sector are higher than normal. We continue to assume that workers must decide which sector to be in before knowing whether there will be a pandemic. Finally, to maintain the symmetry across sectors (which makes the implications particularly clear), we assume that in a pandemic each sector is equally likely to be the one that shuts and each is equally likely to be the essential high-risk one.

In the absence of incentive or fairness considerations, the efficient policy with a high-risk sector is very similar to that of the baseline case. Workers are equally divided among the sectors. In a pandemic, each individual's consumption of the outputs of the two sectors that stay open is unchanged (while their consumption of the output of the sector that shuts of course falls to zero). As before, the efficient allocation can be obtained by a social insurance program of targeted transfers to the unemployed.

With this policy, however, workers in the high-risk sector are worst off, individuals who had been working in the sector that shuts are best off, and workers in the sector that stays open but isn't high-risk are in between. As a result, the presence of an essential high-risk sector has more interesting implications if allocations that make the employed worse off than the unemployed aren't possible. With this constraint, as before, workers in the sector that stays open but isn't hazardous must have higher consumption than the unemployed. But in addition, because working in the high-risk sector involves greater disutility than working in the other open sector, the consumption of workers in the high-risk sector needs to be even greater than that of workers in the other open sector.

This allocation can't be implemented just by targeted transfers to the unemployed; there must also be transfers to workers in the high-risk sector. This implication corresponds to the idea of "hazard pay": workers in the sector that is essential but hazardous are compensated for the additional risks they face by continuing to work.

### Heterogeneous Incomes and Self-Insurance

In the baseline case, individuals in the sector that shuts down are fully compensated for the loss of the income they'd use to purchase the output of the other sector. It follows that if earnings are heterogeneous, individuals with much higher non-pandemic earnings who lose their jobs receive much larger social insurance payments than lower-earning individuals. And although the situation with incentive or fairness considerations is somewhat more complicated, the result that optimal social insurance payments rise strongly with non-pandemic earnings carries over to that case.

This implication of our framework doesn't correspond well with intuition, or with what is done in almost all social insurance programs. For example, although unemployment insurance payments are an increasing function of prior earnings, they're normally capped at a relatively low level, and that policy continued during the pandemic. Similarly, the various rounds of direct payments to individuals excluded those with high incomes, and there were limits on the amount of an individual's earnings that could be paid using funds from the Paycheck Protection Program (where the government made forgivable loans to small businesses hurt by the pandemic).

This gap between our framework's implications and actual policy reflects, at least in part, the framework's omission of two important factors. One is that insurance (whether provided privately or through the government) involves administrative and related costs, which reduces the optimal amount of insurance. The other is that individuals make decisions over multiple periods. As a result, they have some ability to self-insure through their saving and borrowing, and this ability is almost certainly greater for individuals with higher incomes.

The availability of self-insurance reduces the amount of insurance the government should provide when that insurance is costly, and the reduction is greater when an individual's ability to self-insure is greater. Specifically, suppose (realistically) that higher-income individuals are more able to self-insure, in the sense that their consumption falls by a smaller proportion if they become unemployed in a pandemic and there's no social insurance. Then, as we show in the appendix, under reasonable assumptions pandemic social insurance payments rise less than proportionately with individuals' normal incomes. And if self-insurance is sufficiently strong at high enough incomes, optimal pandemic social insurance doesn't provide very high income individuals with any payments at all.

### The Possibility of an Aggregate Demand Shortfall and a Need for Stimulus

A pandemic is a disruption to the economy's ability to produce, so it's natural to describe it as a shock to aggregate supply. But as argued informally by Rowe (2020) early in the pandemic and formalized soon after by Guerrieri et al. (2020), a pandemic that disproportionately affects some sectors can lead endogenously to a fall in aggregate demand larger than the fall in aggregate supply. This result holds in a natural multi-period version of our framework.

Following Guerrieri et al., we say aggregate demand falls by more than aggregate supply in a pandemic if with the previous real interest rate and the optimal social insurance policy, demand for output is less than the efficient level of output (which is of course less than normal output because of the need to shut down some of the economy). If this occurs, obtaining the efficient level of output requires some source of additional aggregate demand, which could come either from a reduction in the real interest rate or from fiscal stimulus.

Whether a pandemic reduces demand by more than supply isn't immediately clear. On the one hand, everyone's consumption falls in a pandemic. In a multi-period setting, this effect tends to make individuals want to borrow against their future income to smooth their consumption. This acts to raise demand relative to supply, and so works in the direction of raising the equilibrium real interest rate. On the other hand, the range of outputs that are available to consume falls in a pandemic because of the shutdowns. This makes consumption less attractive than usual, which works in the other direction.

It turns out that in the most natural multi-period extension of our baseline case, these two effects exactly balance. With optimal social insurance in the baseline case, each individual's consumption of the output of the sector that stays open is the same as in normal times. Together with the assumption that consumption of one sector's output doesn't affect utility from consuming the other sector's, this implies that each individual's marginal utility is the same as normal. As a result, in a multi-period setting, the same real interest rate that makes individuals want to buy the usual level of output in normal times makes them want to buy the efficient level of output in a pandemic. Thus, there's no force causing the equilibrium real interest rate to either rise or fall in a pandemic.

Crucially, however, this result fails if social insurance doesn't equalize consumption across individuals in a pandemic (as occurs with incentive or fairness considerations or with costly insurance). In this case, since total output of the sector that stays open is the same as normal, in the optimal allocation some individuals consume more of the sector's output than normal and some consume less. Those consuming more want to save, which reduces demand, while those consuming less want to borrow, which raises demand. If there are smaller barriers to saving than to borrowing, which is surely realistic, the saving effect dominates the borrowing effect, and so there's an aggregate demand shortfall. Thus, a pandemic is likely to lead to a need for aggregate demand stimulus.

Guerrieri et al. (2020) explore these issues in much more depth. For our purposes, the key message is that a need for stimulus in a pandemic isn't just possible but plausible despite the fall in the economy's safe capacity. One implication is that in thinking about possible social insurance policies in a pandemic, it's appropriate to be concerned about whether they raise aggregate demand in addition to providing insurance.[4]

[4] The appendix discusses two issues concerning pandemic fiscal policy that fit in the social insurance framework but that we don't consider in the text: "forgiveness" (that is, policies that eliminate either payment obligations such as student loan payments, or outstanding debts such as student loan balances),

## Implications for Unemployment Insurance in the Pandemic

A key implication of the social insurance framework is that benefits should be targeted to those who suffer direct economic harms from the pandemic. Because being unable to work is a powerful indicator of being harmed economically, it follows that unemployment insurance should be a central component of pandemic fiscal policy. We therefore turn to what the social insurance perspective implies about how unemployment insurance should be structured in a pandemic, and compare those prescriptions with what was actually done in the United States.

### Coverage

A first issue is who should be eligible for unemployment insurance. In practice, eligibility in normal times is restricted in various ways. Among the most important are that workers must have earned at least a certain amount over some length of time before becoming unemployed (with earnings from self-employment not counting toward this requirement); they must not have left their previous job voluntarily; and they must be either actively seeking work or awaiting recall to their previous job. Yet if individuals could purchase pandemic insurance, they would want to insure themselves against being unable to work in a pandemic regardless of whether they satisfied these criteria. Thus, the social insurance framework implies that unemployment insurance should be available broadly to individuals who aren't working because of a pandemic.

Actual policy in the United States largely followed this implication of our framework. In response to the pandemic, policymakers increased the length of time individuals could receive benefits and relaxed search requirements, waiting periods, and some eligibility rules for regular unemployment insurance. They also enacted the Pandemic Unemployment Assistance program, which extended coverage to many previously uncovered workers, including the self-employed, gig workers, and workers who quit their jobs because of childcare needs resulting from the pandemic

---

and aid to state and local governments. It also describes some extensions that have little effect on the main messages of our analysis.

(Cajner et al., 2020). Finally, as discussed in more detail below, they added a fixed weekly amount to benefits (including those through Pandemic Unemployment Assistance) for much of the pandemic.

As Cajner et al. (2020) stress, a host of factors—ranging from fraudulent claims to idiosyncratic biweekly reporting in some large states—make it very challenging to determine the number of individuals receiving benefits during the pandemic. Nonetheless, it's clear that coverage rose broadly in line with job losses. Employment (as measured by the household survey, and so including the self-employed) fell by 24.7 million from February to April 2020.[5] Cajner et al.'s preferred measure of continued regular unemployment insurance claims rose by 19.7 million from its prepandemic level to its peak. Determining the number legitimately receiving Pandemic Unemployment Assistance benefits is even harder, but Cajner et al.'s results indicate it was at least several million.

Of course, coverage wasn't complete. Some of the pandemic unemployed, such as most new entrants to the labor force, weren't eligible for any unemployment insurance. And, the conjunction of the enormous crush of claims, limited processing resources, and antiquated administrative systems in many states led to long delays in processing claims (Bitler, Hoynes, and Schanzenbach, 2020; Cajner et al., 2020). Nonetheless, there was unusually broad coverage of the unemployed in the pandemic.

## Duration

A second issue that's not directly addressed by the one-period framework is how long pandemic-related unemployment insurance should last. Again, however, the logic of the social insurance perspective is clear. If individuals could purchase pandemic insurance, they would want to be insured for the duration of the pandemic. Thus pandemic-related unemployment insurance

---

[5] Because the impact of the pandemic swamped the effects of normal seasonal variation, the numbers in this paragraph are not seasonally adjusted.

should be available for however long the pandemic lasts.

Actual U.S. policy departed considerably from this implication of the social insurance framework. Policymakers did increase the number of weeks that individuals could receive benefits. But throughout the pandemic, the various extensions of duration, expansions of coverage, and increases in benefits were tied to calendar time rather than to metrics of the course of the pandemic or the state of the economy. Moreover, there was often great uncertainty about whether the emergency measures would be extended or allowed to lapse, which is not what beneficiaries would desire from a social insurance program.

Consider, for example, the Pandemic Unemployment Assistance program. The original Coronavirus Aid, Relief, and Economic Security (CARES) Act, enacted on March 27, 2020, provided 39 weeks of benefits under the program and set its end data as December 31, 2020. The program was continued at the last minute through mid-March 2021 by the Consolidated Appropriations Act, 2021, enacted December 27, 2020, which also provided an additional eleven weeks of benefits. The program was extended again (through September 6, 2021) by the American Rescue Plan Act of 2021, enacted on March 11, 2021—just a few days before the program was to end.

Similarly, the additions to benefits fluctuated greatly in ways largely unrelated to the severity of the pandemic. The CARES Act provided an extra $600 per week of unemployment insurance benefits through July 31, 2020, at which point the supplement was allowed to lapse. The Lost Wages Assistance Program through the Federal Emergency Management Agency then provided $300 per week of additional benefits (or, in a few states, $400 per week) for roughly six weeks for many workers, though workers receiving low benefits were excluded. Additional benefits then lapsed entirely through the end of 2020. The Consolidated Appropriations Act, 2021 provided $300 per week in additional benefits from January 1, 2021 through March 14, 2021. Finally, the American Rescue Plan Act of 2021 extended the $300 per week through September 6, 2021.

**Replacement Rates**

An issue for which it's much harder to determine the exact implications of the social insurance framework is how generous unemployment benefits in a pandemic should be, and how that generosity should vary with prior income. In our baseline case, individuals who remain employed are taxed to finance transfers to the unemployed, and insurance payments to the unemployed are less than their prior incomes. In practice, however, additional unemployment insurance in the pandemic has been financed in the short term by government borrowing. In this case, as we describe in the appendix, the natural modification of our baseline case implies that optimal social insurance fully replaces the incomes of individuals who can't work because of the pandemic.

The various extensions of our baseline case discussed in the previous section, however, suggest two reasons for less than 100 percent replacement. First, there may be incentive or fairness considerations that imply that the unemployed shouldn't be better off than the employed. Since the unemployed don't have the disutility of working, in a one-period setting this requires that they have lower incomes than the employed. As various authors have pointed out, however, things are more complicated in a dynamic setting. Someone who's unemployed may face a long and challenging job search when the pandemic ends. Thus it's not clear that replacement rates need to be much (or even at all) below 100 percent during the pandemic to prevent individuals preferring unemployment to employment (for example, Ganong, Noel, and Vavra, 2020).

Second, and more importantly, providing social insurance involves costs beyond the insurance payments. Those costs include administrative and related expenses, the distortions caused by the incentive effects of the insurance, and the distortions from raising the additional revenues at some point to satisfy the government's intertemporal budget constraint. These costs imply that even in the absence of incentive or fairness issues, optimal social insurance should leave the consumption of the unemployed somewhat below that of the employed. In the appendix, we argue that 10 to 15 percent is a plausible ballpark figure for the size of the shortfall.

How much social insurance is needed to achieve this level of consumption for unemployed workers depends crucially on the possibility and cost of self-insurance. In particular, to make more precise statements about optimal replacement rates and how they should vary with prior income, we would want two types of information. The first is a comparison of the consumption of unemployed and otherwise similar employed individuals in the absence of unemployment insurance, as a function of their income. This would be informative about the extent of self-insurance. The second is evidence about how the unemployed finance their consumption in the absence of unemployment insurance. This would be informative about the cost of self-insurance. To the extent individuals use types of self-insurance more costly than government-provided social insurance (as might be true of borrowing at very high interest rates or selling highly illiquid possessions), it's optimal for the government to provide the insurance. On the other hand, if sufficiently high-income individuals can do enough self-insurance at a cost less than that of the government to keep their consumption from falling very much if they become unemployed, optimal social insurance doesn't provide them with any social insurance payments.

Ganong and Noel (2019) examine the consumption behavior of unemployment insurance recipients before the pandemic. In a sample where recipients are probably financially healthier than typical beneficiaries, they find that after recipients exhaust their benefits, their consumption is on average 20 to 30 percent below what it was when they were employed. They also find that the falls average 30 to 40 percent among those in the highest third of the ratio of benefits to prior income and in the lowest third of assets relative to prior spending (both which tend to be associated with lower incomes). Farrell et al. (2020) and Ganong et al. (2021a) find that delays in benefits early in the pandemic also resulted in large reductions in consumption for the unemployed.

Another type of evidence about these issues comes from short-run marginal propensities to consume (MPCs) out of stimulus payments during the pandemic. A high MPC is suggestive of current marginal utility being high relative to future marginal utility, and thus of short-run financial distress. Karger and Rajan (2021) and Baker et al. (2021) find large MPCs from the stimulus payments. More importantly for our purposes, they find that it varied substantially with income. For example, Baker et al.'s point estimates (from their Table 5, Column 1) imply an MPC over 3 months of 0.66 at a monthly income (net of withholding) of $500 and just 0.15 at a monthly income of $5000. This evidence points to greater economic stress during the pandemic among low-income individuals.

Taken together, these studies indicate that unemployed workers may have limited ability to self-insure, and that this ability may be substantially smaller among lower-income workers. This suggests that the replacement rate for unemployment benefits may need to be fairly substantial, though clearly less than 100 percent, to result in a loss of consumption in the 10 to 15 percent range. It also suggests that replacement rates should decline as prior income rises. However, the existing evidence is not enough to pin down optimal replacement rates precisely.

Even though we are unable to say what exactly replacement rates from unemployment insurance during a pandemic should be, it is clear that actual replacement rates have differed sharply from the prescriptions of a social insurance perspective. Ganong, Noel, and Vavra (2020) show that the flat $600 per week of additional benefits raised replacement rates to well over 100 percent for most workers.

There appear to have been two forces behind the policies involving greater than 100 percent replacement. One is the pursuit of other objectives, especially redistribution toward lower-wage workers and aggregate demand stimulus. The other is idiosyncratic factors: Ganong, Noel, and Vavra (2020) report that an overestimate of the average wage of workers who would lose their jobs led policymakers to underestimate the impact of the $600 weekly adjustment on replacement rates, and that the very limited capacities of state unemployment insurance systems led policymakers to adopt the fixed supplement rather than more complicated additions to benefits.

The later iterations of benefit increases, which were in the range of $300 per week, likely reduced replacement rate to below 100 percent for many workers, and so are more consistent with the social insurance perspective. And, regular unemployment benefits continued to phase out at high incomes, which is consistent with the social insurance framework given the evidence that self-insurance is easier for high-income unemployed workers.

**Unemployment Insurance Payments as Demand Stimulus**

As we have discussed, even with optimal social insurance, a pandemic can lead to a shortfall of aggregate demand from what's needed to yield the desirable level of output. Thus, even though the central motivation for broad unemployment insurance in a pandemic is social insurance, it is valuable to know whether unemployment insurance is also an effective source of aggregate demand stimulus.

Three types of evidence suggest that it is. First, the social insurance framework implies that unemployment insurance should be targeted to individuals whose marginal utility from consumption would otherwise be temporarily high. Such individuals would be expected to devote a large fraction of a marginal increase in resources to current consumption.

Second, the evidence from before the pandemic points to a high MPC out of unemployment insurance benefits, and thus to it being effective stimulus. Ganong and Noel (2019) report a very conservative lower bound for the MPC out of unemployment insurance benefits of 0.27, an upper bound of 0.83, and a point estimate under reasonable assumptions of 0.77. In a survey of work in this area, Chodorow-Reich and Coglianese (2019) suggest an MPC of about 0.35 over one month and 0.55 over a year.

Third, the evidence from the pandemic also points to high MPCs out of unemployment insurance benefits. Most notably, Ganong et al. (2021a) estimate an MPC out of benefits (both regular and the $600 supplement) between 0.29 and 0.43 over one month, and between 0.62 and 0.69 over six months.

One wouldn't expect the large MPCs to apply regardless of the level of benefits, however. If benefits are sufficiently high that the marginal utility of current consumption for the unemployed is driven down to its normal level, a large fraction of any additional benefits is likely to be saved. The evidence doesn't clearly support this prediction, however. Ganong et al. (2021a) find that because of the high replacement rates, the unemployed had both unusually high consumption and large increases in their stock of savings during the period when benefits included the $600 supplement. Furthermore, despite the higher saving, they estimate an MPC out of the subsequent $300 supplement that is very similar to the MPC out of the initial benefits. Thus, the issue of under what circumstances MPCs are high remains open.

## Implications for Hazard Pay in the Pandemic

As discussed in the first section, a novel implication of the social insurance perspective is that in some cases the government should provide hazard pay to high-risk essential workers during a pandemic. This implication results from extending the baseline model to include a third sector where workers doing very socially valuable jobs face greater risk of illness, and hence have greater disutility of work during a pandemic. This, combined with either notions of fairness or difficulty in determining who should be working, leads to paying such at-risk workers a bonus. Intuitively, the need for hazard pay comes from seeking to make sure workers employed in the riskier sector aren't worse off than workers whose sector had to shut down. This section investigates this implication of the social insurance perspective on fiscal policy in more detail. It also considers some of the many practical issues involved in the design and implementation of an actual program.

Government-funded hazard pay was considered during the pandemic, and actually implemented on a very limited basis. These programs provide a useful baseline for discussion. The 2020 HEROES Act, which passed by the House of Representatives in May 2020, contained a $200 billion hazard pay program (U.S. Congress, 2020). The proposed program, which did not make it into final legislation, would have paid an extra $13 per hour for the period from January 27, 2020 (the first day of the declared public health emergency) until 60 days after the last day of the emergency. The definition of who was potentially eligible was very broad. Qualifying work was defined as work "not performed while teleworking from a residence," and involving "regular in-person interactions with … patients; … the public; … or coworkers" in a number of specific areas (116th Congress, 2020). Actual eligibility determination required an application by the employer to the Treasury Department. If the application was approved, the employer would then receive a grant and be responsible for providing the premium pay to workers.

An example of a hazard pay program that was actually implemented is the Covid-19 Pennsylvania Hazard Pay Grant (Pennsylvania Department of Community and Economic Development, 2020), enacted in May 2020. The employees receiving the hazard pay needed to interact with others, and couldn't be teleworking from home. The list of eligible industries was decidedly shorter than the HEROES Act program. The amount of hazard pay was $3 per hour, up to a total of $1200. Like the HEROES Act, the Pennsylvania program required employers to apply for the grants. Employers could only apply for grants covering up to 500 full-time-equivalent employees per location. The budget for the program was $50 million.

**Who Should Receive Hazard Pay?**

Whether some workers should receive hazard pay depends crucially on differences in risk across occupations during the pandemic. A number of studies have tried to investigate this issue. Their findings are summarized in Table 2.

One of the most detailed studies combines comprehensive health records and occupation data to form a complete individual-level national database for Norway (Magnusson et al., 2021). The study then identified occupations with direct contact with children, students, patients, and customers, and compared Covid-19 infection rates for these occupations with those of all other working-age adults. The study found that after controlling for age and sex, healthcare workers and transit workers had approximately three times the risk of contracting Covid-19 during the first wave of infection (February 26–July 17, 2020) (Magnusson et al., 2021, p. 7). Interestingly, in the second wave (July 18–December 18, 2020), healthcare workers no longer had significantly elevated Covid risk. Instead, the occupations with the highest odds ratios of infection were food service workers (food counter attendants, bartenders, and waiters), transit workers, and cleaners (p. 8).

A study of 120,000 people in the United Kingdom compared the risk of Covid-19 (severe enough to be diagnosed in a hospital or emergency room) between workers in eight essential occupations and nonessential workers (Mutambudzi et al., 2021). The highest relative risk was for healthcare professionals, medical support staff, health associate professionals, and social care workers. Healthcare workers had more than a seven-fold increase in Covid-19 risk. A strong point of this study is that it was able to control for a wide range of health and demographic covariates.

Song et al. (2021) used detailed private insurance data from Pennsylvania to compare Covid-19 inflection rates for essential and nonessential workers. Essential workers were defined as those employed by firms designated by the governor as life-sustaining businesses and permitted to continue physical operations during a statewide shutdown. Because this definition required physical operation, many of the workers involved were likely to have contact with others. The authors use a difference-in-differences specification to see if essential workers had higher infection rates following the shutdown order. They find that for a sample of policyholders younger than 65, “being an essential worker is associated with a 53% increase in likelihood of being Covid-positive” (p. 11). This increase was most notable in the health and social care sectors, but still present to a smaller extent for essential workers outside those sectors.

Two other studies use death records to assess the occupational mortality risk of Covid-19. A Swedish study found that taxi and bus drivers and service sector workers had a higher risk of dying from Covid-19 than other workers, but the elevated risk disappeared when demographic covariates were included (Billingsley et al., 2020). A California study estimated excess deaths for nine occupational sectors over the period March to November 2020 (Chen et al., 2021). The sectoral highs were in food or agriculture, transportation or logistics, manufacturing, and facilities. Healthcare workers faced only slightly elevated mortality risk. A weakness of this study is that the authors have little information on covariates, as well as imperfect data on occupation.

Based on this evidence, it appears that hazard pay would have made sense during the pandemic. There was indeed a portion of the economy that needed to remain open but that involved a greater risk of illness. At the same time, the evidence suggests that the fraction of essential workers at noticeably higher risk was relatively small. Both Magnusson et al. (2021) and Mutambudzi et al. (2021) identify healthcare workers as being the main workers at greater risk (at least early in the pandemic). Using occupational data from the May 2019 National Occupational Employment and Wage Estimates, healthcare workers (broadly defined) account for about 10 percent of all U.S. workers (U.S. Bureau of Labor Statistics, 2019). Expanding the list to include food service workers, taxi drivers, and cleaners, the other noticeably higher risk occupations, adds another 8 percent of U.S. workers.

Compared with this list, the HEROES Act had an excessively broad definition of who should receive hazard pay during the pandemic. The list of eligible occupations or industries in the bill went on for more than four pages and included not just the obvious ones of healthcare and public transportation, but child care, barge operations, longshoremen, laundry work, and restaurant work. The list of eligible industries in the Pennsylvania program was shorter, but still included many for which there's little evidence of higher risk, such as security services, freight trucking, and retail food stores. Both programs did limit the bonuses to frontline workers—that is, essential workers who were required to interact with others. But, Blau, Koebe, and Meyerhoff (2021) estimate that in March 2020, about 43 percent of all U.S. workers met these criteria.

Both the HEROES Act and the Pennsylvania hazard pay program relied on firms to apply for hazard pay funds from the government. From a social insurance perspective, this feature is problematic. Relying on employer application risks identical workers being treated differently (or worse, workers facing less risk getting premium pay and workers facing more risk not). Whether workers receive the government-provided hazard pay could depend on employer motivation or simply on how quickly employers respond to the call for applications. These outcomes aren't socially optimal.

Various alternatives are possible. A minor variation that would reduce the possible randomness would be for the government to invite firms in the relevant industries to say how many workers they have eligible for the hazard pay. The presumption would be that all such identified workers would receive it. Firms would still have to make some effort for their workers to receive the bonus, but the burden would be less. A more extreme variation would be to require individual workers to apply, as they do for unemployment insurance. This proposal has the great disadvantage of being administratively complex. Not only would the government need to process a very large number of applications, but firms would likely still have to be involved because they are the sensible unit for dispensing the pay.

### How Large Should the Hazard Premium Be?

The HEROES Act called for a hazard premium of $13 per hour for all eligible workers; the Pennsylvania program provided a boost of $3 per hour. Since risk is related to contact, and contact is time-dependent, it is appropriate for the hazard premium to be per hour worked, rather than lump sum. A more difficult and fundamental question is just how much extra pay would be appropriate for high-risk workers.

A crude way to try to answer this question is to blend risk estimates with an estimate of the value of a statistical life. Table 3 presents such a calculation. The Centers for Disease Control and Prevention (CDC) provides estimates of the number of Covid-19 infections for adults 18 to 64 in the United States over roughly the first year of the pandemic (CDC, 2021a). We combine this with data on the working-age population in 2019 from the U.S. Census Bureau (2020) to get an overall infection rate (row 1 of Table 3). Magnusson et al. (2021, Figure 3) find that healthcare and public transit workers in Norway had a risk of contracting Covid-19 that was approximately three times greater than other working-age adults (adjusting for age and sex) during the first wave of the pandemic. As discussed above, these workers comprise about 10 percent of employed workers in the United States, or 7 percent of the working-age population in 2019. From this, we can back out estimates of the risk of getting Covid-19 during the first year of the pandemic for both the high-risk group and other working-age adults (row 2). This calculation suggests that high-risk workers had a 59 percentage point higher risk of infection.[6] We then estimate the infection fatality rate for working-age adults by dividing confirmed working-age deaths (CDC, 2021b) by estimated cases (row 3).[7] The EPA and other government agencies use $10 million as their estimate of the value of a statistical life (row 4) (see, for example, U.S. Environmental Protection Agency, 2016).

Armed with these components, we can calculate the cost to high-risk workers of the extra risk of dying they incurred during the pandemic. Multiplying the excess risk times the infection fatality rate times $10 million yields $8617 (row 5 of Table 3). The pandemic had been going on for roughly a year as of the time of the data used in the calculation, so this estimate corresponds to the pay boost that would be needed for a year to compensate for the additional risk. This works out to $4.13 per hour for a typical 2087-hour work year (U.S. Office of Personnel Management, 2021). This is roughly what the Pennsylvania hazard pay program provided, but substantially less than what the HEROES Act called for.

There are various ways that this estimate of the appropriate size of the hazard premium could be too low. Most obviously, it only includes the risk of dying from Covid-19. Studies suggest that at least 10 percent of Covid survivors experience long-term effects (see, for example, Greenhalgh et al., 2020). Workers might put a substantial dollar value on that possibility as well. Second, hazard pay is motivated by a desire to ensure that high-risk workers remain willing to do socially valuable work. As a result, what matters is not necessarily their actual risk of infection and death,

---

[6] This difference is almost surely somewhat too large. It may reflect the fact that the Magnusson et al. (2021) estimates for relative risk come from early in the pandemic when overall infection rates in Norway were very low. It's worth reiterating, however, that Mutambudzi et al. (2021) also estimate very high risk for healthcare workers.

[7] Our resulting infection fatality rate of 0.146 percent matches quite closely the estimates from O'Driscoll et al. (2020) for the infection fatality rate for working-age adults based on a meta-analysis of studies for a number of countries.

but their perceived risks. If high-risk workers systematically overestimate their risks, the hazard premium might need to be higher. For example, if the perceived infection fatality rate were twice as high as the CDC estimates, then the perceived cost of excess risk to high-risk healthcare and transit workers would be over $8 per hour—getting closer to the hazard premium called for in the HEROES Act.

Importantly, the calculation above is only for the highest-risk workers. The HEROES Act proposed to provide hazard pay for most frontline workers. The high-contact workers analyzed by Magnusson et al. (2021) account for roughly 31 percent of all workers in the United States in 2019.[8] For this broader group, the overall weighted average of the added risk of infection (adjusted for age and sex) was a factor of 1.59 (which is very similar to the factor of 1.53 found by Song et al., 2021). Line 6 of Table 3 shows that the dollar value of the extra risk to frontline workers as a group is relatively low—precisely because most face little elevated risk (or in some cases lower risk than non-frontline workers). The calculation suggests that the dollar value of the added risk for this broad group is $2704, or roughly $1.30 per hour.

Given the variation in risk among frontline workers, an appealing possibility would be to have some sort of tiered hazard pay system. While neither the evidence nor administrative capacity is likely to be adequate for a highly variable scale, it seems possible that having a few tiers could result in a program that more closely fits the social insurance ideal. For example, in the Magnusson et al. (2021) study, roughly 10 percent of workers have a threefold increase in risk, and another 8 percent have a twofold increase.[9] As shown in line 7 of Table 3, these numbers imply a dollar value of the extra risk of $4.13 per hour for the highest-risk workers, and $2.06 for the second highest risk category.

While these various calculations do not yield a firm number for the size of any hazard

[8] We calculate this percentage by matching the categories of workers in the Magnusson et al. study with those from the Bureau of Labor Statistics (U.S. Bureau of Labor Statistics, 2019).
[9] For the increased risk for food service workers, taxi drivers, and cleaners, we use the estimates from the second wave of infection, when the relevant sectors were more likely to be open (see Magnusson et al., 2021, Figure 4).

premium, they illustrate how policymakers might go about figuring out an appropriate premium in a pandemic. A crucial input, particularly for a tiered hazard pay scheme, is good evidence on the degree of elevated risk for different occupations and industries.

There are obviously other ways policymakers could try to ascertain the appropriate amount of hazard pay needed. For example, they could gather data on labor shortages in frontline industries. During the pandemic, there was a great deal of anecdotal evidence of labor shortages, particularly for nurses and home healthcare workers (see, for example, McLernon, 2020, and Nguyen, 2020). Though it would be impractical to try to estimate a full labor supply function in real time, the degree of labor shortage in frontline industries could provide a rough guide for at least the appropriate ranking of the hazard premium across industries.

**Should Hazard Pay Phase Out at High Incomes?**

Whether hazard pay should phase out above some income level depends on the relationship between the utility of income from working and the disutility of working itself for high-wage, high-risk workers. If the two were reasonably close before the pandemic, the rise in the disutility of work caused by the pandemic could cause many of these workers to stop working (even if they were not eligible for unemployment insurance). Since such workers are very valuable to society, in this case hazard pay should continue even at high incomes. On the other hand, if the utility provided by the income from working was substantially above the disutility of working for high-income workers in the high-risk sector before the pandemic, even a substantial rise in their disutility of work would not make them choose unemployment. In this case, it would make sense for hazard pay to phase out.

There are good reasons for thinking that the scenario where high-income workers in the risky sector choose to continue working despite increased risk is the more plausible one. Workers such as doctors and dentists have invested years in training and building a practice. The present value of the income loss they would face from leaving their jobs, even temporarily, is very large. More generally, it is likely that high human capital and high match quality create a substantial wedge between what high-skilled, high-risk workers earn and the disutility of work. Finally, there's also evidence that nonpecuniary job benefits are higher for well educated workers (see Duncan, 1977). This too creates a wedge between the wage and the disutility of labor that would encourage high-income, high-risk workers to continue working even if the disutility of work rises. Thus, it would likely be appropriate for hazard pay to phase out at high incomes.

Interestingly, the HEROES Act did not phase out hazard pay entirely for high-wage workers, though it did limit it. The act capped the amount of premium pay at $10,000 for eligible workers earning less than $200,000, and $5000 for eligible workers earning more than $200,000. The Pennsylvania program, on the other hand, only applied to workers earning less than $20 per hour (excluding benefits and overtime).

**Hazard Pay as Demand Stimulus**

The primary motivation for hazard pay during a pandemic from a social insurance perspective is to make sure essential workers in hazardous sectors aren't worse off than workers who become unemployed because their sectors shut. However, because there may be inadequate aggregate demand during a pandemic, an important practical consideration is whether hazard pay is likely to be useful as demand stimulus. From a logical standpoint, hazard pay wouldn't seem to have a particularly large bang for the buck. By definition, the workers receiving it remain employed. As a result, one might think that the recipients have a lower marginal propensity to consume than the unemployed workers discussed in the previous section.

One factor that militates against this presumption is that many frontline workers are relatively low-income. Blau, Koebe, and Meyerhofer (2021, p. 172) estimate that frontline workers had lower average hourly wages ($22.76 versus $27.05 for all workers). This is even true for healthcare workers, where the case for hazard pay is strongest. While healthcare practitioners and technicians earn more than $40 per hour, Blau, Koebe, and Meyerhofer find that healthcare support workers earn less than $20 per hour (2021, p. 172). The other five lowest-wage frontline occupations the authors identify are food preparation and serving, building and grounds cleaning and maintenance, personal care and service occupations, farming, fishing, and forestry, and transportation and material moving. Many of these are occupations that would be next in line after healthcare to receive hazard pay based on their relative Covid-19 risk.

Other studies reach similar conclusions about the low-wage nature of much elevated-risk work. Kinder, Stateler, and Du (2020) calculate that “as of 2018, nearly half (47%) of all frontline essential workers earned less than a living wage” (which they identify as $16.14 in 2018). For example, personal care aides had a median wage of $11.55, and janitors and cleaners had a median wage of $12.55. Kearney and Muñana (2020), using the Kaiser Family Foundation Tracking Poll, find that essential workers working outside the home were more likely to earn less than $40,000 than other currently employed workers (31 percent versus 19 percent). 49 percent of them said they would struggle to pay a $500 unexpected medical bill, while only 31 percent of other workers said they would.

That high-risk workers are disproportionately low-wage and would struggle to pay bills suggests that they are likely to have a higher MPC than a typical worker. A higher MPC should translate into a higher fiscal multiplier. Thus, hazard pay should pack a larger stimulatory impact than broader types of stimulus, such as general tax cuts or widely available one-time payments. This extra stimulatory impact could be enhanced by limiting hazard pay to low-wage workers, as in the Pennsylvania program.

## Conclusion

In many ways, our social insurance perspective on fiscal policy harkens back to an older literature. Rather than focusing on aggregate demand management, we emphasize the more traditional role of government in providing insurance against life’s vicissitudes. This is particularly appropriate for fiscal policy during a pandemic, when output needs to remain low in some sectors for health reasons, and aggregate demand stimulus cannot flow to many affected workers because their sectors are shut.

Our framework shows that thinking in terms of social insurance leads naturally to directing government aid to those directly harmed by the pandemic—particularly the unemployed and those who work in essential jobs with a high risk of infection. Our more practical analysis of unemployment insurance and hazard pay shows how a social insurance perspective can provide guidance on who should receive these benefits, how large they should be, and how long they should last. An analysis of the policies actually taken or proposed for unemployment insurance and hazard pay during the pandemic shows that many of the actions follow what a social insurance perspective would recommend, but a number did not.

An obvious question about the social insurance perspective on fiscal policy is whether it's likely to have usefulness beyond the Covid-19 pandemic. We believe it does. First, it is all too likely that there will be future pandemics. The implications of the social insurance perspective we have discussed could yield much more successful fiscal policy in any future public health crises.

Second, it's possible that the United States and other countries could face more regionally-concentrated recessions in the future that share important similarities to a pandemic. For example, parts of the American West and of Australia have experienced prolonged economic disruption due to drought-induced fires. As with a pandemic, such natural-disaster-fueled downturns aren't easily remedied with broad aggregate demand stimulus. They're also likely to involve high risks to certain essential workers. The lessons derived from a social insurance perspective for pandemics about targeting relief to those most affected and paying a hazard premium to essential high-risk workers are likely to carry over to these downturns as well.

Finally, while the unique features the pandemic have made a social insurance perspective vital in the current situation, it may have value in more ordinary downturns. Most recessions involve highly unequal impacts on different types of workers, and general stimulus often takes a long time to help some workers regain employment. A fiscal response that focuses on both aggregate demand management and social insurance might prove more effective in dealing with future recessions, regardless of their cause.

**Table 1**

**Variations on the Baseline Case**

| Variant | Key Implication |
|---|---|
| 1. Incentive or fairness considerations | Unemployed workers have lower consumption than the employed |
| 2. A high-risk essential sector (plus incentive or fairness considerations) | Hazard pay for high-risk essential workers |
| 3. Heterogeneous incomes, possibilities for self-insurance, and costs of social insurance | Higher replacement rates of lost income for lower-income workers |
| 4. Multiple periods, barriers to borrowing, and a force causing the unemployed to have lower consumption than the employed | Aggregate demand falls by more than aggregate supply, implying a role for aggregate demand stimulus |

**Table 2**

**Summary of Studies of Occupational Risk of Covid-19**

| Study | Sample | Riskiest Occupations | Degree of Elevated Risk |
|---|---|---|---|
| | | ***Outcome Variable: Covid-19 Infection*** | |
| 1. Magnusson et al. (2021) | All working-age Norwegians | 1st wave: healthcare, public transit workers<br>2nd wave: food service, public transit workers, cleaners | Approximately 3 times all other working-age adults<br>Approximately 1.5–2 times all other working-age adults |
| 2. Mutambudzi et al. (2021) | 120,075 UK Biobank participants | Healthcare workers<br>Social care workers | 7.43-fold increase in risk<br>2.46-fold increase in risk |
| 3. Song et al. (2021) | 400,000 primary policy holders of a major insurer in PA | Essential workers as revealed by whether their industry was shut down | 53% increase in risk relative to non-essential workers |
| | | ***Outcome Variable: Covid-19 Death*** | |
| 4. Billingsley et al. (2020) | All deaths in Sweden 3/5 to 5/7/2020 | Taxi and bus drivers<br>Service sector | 3.7 times the risk of IT techs.<br>2.2 times the risk of IT techs. (not statistically significant) |
| 5. Chen et al. (2021) | Deaths in CA from 1/2016 to 11/2020 of people 18–65 | Food or agriculture<br>Transportation or Logistics<br>Manufacturing<br>Facilities | 1.39 times expected deaths<br>1.31 times expected deaths<br>1.24 times expected deaths<br>1.23 times expected deaths |

**Table 3**

**Putting a Dollar Value on Increased Covid-19 Risk for High-Risk Workers**

1. Share of working-age population who had Covid-19 during first year of the pandemic:
   75,302,292/213,610,414 = 0.35

2. Estimated risk of getting Covid-19 based on Magnusson et al. (2021):
   Other working age population: 29%
   Healthcare and transit workers: 88%
   Difference: 59 percentage points

3. Estimated infection fatality rate for working-age population:
   110,143/75,302,292 = 0.00146

4. Value of a statistical life: $10 million

5. Dollar value of excess risk to healthcare and transit workers:
   0.59 • 0.00146 • $10 million = $8617 per year (or $4.13/hour)

6. Dollar value of excess risk to all frontline workers based on Magnusson et al.:
   0.18 • 0.00146 • $10 million = $2704 per year (or $1.30/hour)

7. Dollar value of excess risk to top two risk tiers based on Magnusson et al.:
   *Healthcare and transit workers*:
   0.59 • 0.00146 • $10 million = $8617 per year (or $4.13/hour)
   *Food service workers, taxi drivers, and cleaners*:
   0.29• 0.00146 • $10 million = $4308 per year (or $2.06/hour)

*Note:* Equations may not hold exactly because of rounding.

# APPENDIX

This appendix is a more formal presentation of the material in the section of the paper on "A Social Insurance Perspective." The organization of the first five sections follows that section of the paper. Sections F and G consider the two extensions of the model that we allude to in the paper but don't analyze in the main text—the possibility of various types of debt and payment "forgiveness," and the treatment of state and local governments. Finally, Section H extends the analysis of incentive and fairness considerations in the text and in Section B to the case where the government has some but less than full information about who should optimally be working in a pandemic.

## A. A Baseline Case

***Assumptions*.** The economy lasts for a single period and consists of a continuum of identical individuals with a total mass of 1. Throughout, all markets are perfectly competitive.

There are two sectors, $A$ and $B$. There are three possible states of the world: no pandemic, $A$-sector pandemic, and $B$-sector pandemic. If a sector isn't affected by a pandemic, its production function is linear in employment in the sector: $Q_i = L_i$ ($i = A, B$), where $L_i$ is the amount of labor employed in sector $i$. If a sector is affected by a pandemic, its production is 0. The probabilities of the $A$-pandemic state and the $B$-pandemic state are equal.

Each individual has utility $U(C_A) + U(C_B) - V(L)$, where $C_i$ is their consumption of the output of sector $i$ and $L$ is the amount they work. The functions are assumed to satisfy $U'(\bullet) > 0$, $U''(\bullet) < 0$, and $V(1) > V(0)$. $U(0)$ and $V(0)$ are normalized to 0. Each individual's labor supply can take on only the values 0 and 1. Individuals are mobile between sectors before the state of the world is realized but immobile ex post. We also assume that the utility from consumption is large enough relative to the disutility of working that in the various cases we consider, in equilibrium individuals work unless prevented by a pandemic.[10]

***<u>Equilibrium under different allocation mechanisms</u>*****.** Suppose first that there are no insurance markets. The symmetry of the sectors causes half of individuals to be in each sector. In the no-pandemic state, output of each sector is ½, the symmetry causes the prices of the two outputs to be equal, and competition and the assumption about the production function imply that the wage in each sector equals the price of that sector's output. Diminishing marginal utility causes each individual to consume ½ unit of each output. If there's a pandemic in one sector, individuals in that sector earn no income, so their consumption of both outputs is 0. Each individual in the non-pandemic sector earns an amount equal to the price of that sector's output, and so consumes 1 unit of that sector's output.

This outcome involves an inefficient allocation across states. In the *A*-pandemic state, the marginal utility (from consumption of the output of sector *B*, which is the only sector that can produce) of individuals in the *A* sector is $U'(0)$, while the marginal utility of individuals in the *B* sector is $U'(1)$. In the *B*-pandemic state, the marginal utilities are reversed. Thus there's scope for insurance across states.

If there are complete markets (achieved either through markets for insurance against a pandemic in each sector or markets in all Arrow-Debreu commodities), it's straightforward to check that the allocation of individuals to sectors and the no-pandemic allocation are the same as without insurance, and that there's full risk-sharing: if a pandemic hits a sector, the consumption

[10] The exact condition needed for individuals to work when they're able to depends on the case considered. For example, in the very first case discussed below (the case of no insurance markets in the baseline version of the model), the required condition is just $U(1) > V(1)$—in a pandemic, an individual in the sector that remains open prefers to work and spend all their income on the output of that sector than to not work and have no consumption. (For individuals to work in the no-pandemic state, the required condition is $U(1/2) + U(1/2) > V(1)$. However, this condition follows from $U(1) > V(1)$ and our assumptions about $U(\bullet)$. Thus we only need to assume $U(1) > V(1)$ for individuals to always work when they're able to in the case of no insurance markets.)

The one place where we depart from the assumption that individuals always work if they're able to is when we introduce partial ex post labor mobility. We model this possibility by assuming a heterogeneous cost of switching sectors among workers in the sector that's shut. In this case, some individuals who could conceivably continue to work but face a very high cost of doing so don't work.

of the output of the sector that stays open is ½ for the individuals in both sectors. To see that this is the equilibrium, note that with these outcomes, in any state every individual has the same marginal utility. Moreover, the allocation satisfies the technological constraints, and individuals in each sector have the same expected utility.

Finally, if there are no insurance markets, the government can use taxes and transfers to implement the allocation that would occur if they were present. In the absence of a pandemic, it takes no action; as a result, the allocation without a pandemic is the same as before. But if a pandemic shuts one sector, it taxes the workers in the sector that remains open half their income and transfers the proceeds to the workers in the sector that shuts. Recall that the wage of each worker in the sector that stays open equals the price of that sector's output. Thus each worker's after-tax-and-transfer income is one-half the price of the output of the sector that remains open, and so everyone's consumption is ½.

***Discussion*.** In the text, we highlight several messages of the baseline case: the optimal social insurance policy consists of targeted transfers; it doesn't involve aggregate demand stimulus; it equalizes income across individuals but doesn't fully replace unemployed workers' lost income; and one can think of it as the government taking the amount the employed would normally spend on the output of the sector that shuts and giving the proceeds to the individuals in that sector.[11]

Here, we highlight an additional implication of the baseline case concerning what happens under complete markets: even though the possibility of a pandemic is a systemic risk, the pricing of insurance is actuarially fair. That is, even though a pandemic reduces total output, when

---

[11] As we describe the optimal social insurance policy, not just the transfers but the taxes to finance them are targeted, since they're only levied on the employed. But an equivalent policy is to tax everyone half their normal income and transfer to each unemployed worker the full amount of their normal income (which would mean that the unemployed would have to pay back half of what they receive in transfers as taxes). With this policy, only the transfers are targeted; but as with the policy described in the text, each individual's after-tax-and-transfer income is half their normal income. Or, each individual could be taxed half their actual income; since the unemployed earn no income, the tax would fall only on the employed. Thus, targeted taxes aren't needed to achieve the efficient allocation.

individuals can obtain insurance they can trade off consumption in a pandemic state (of the output that can be produced) and consumption in the no-pandemic state (of either output) just according to the relative probabilities of the two states. The reason is that a pandemic doesn't reduce consumption of both sector's outputs, but instead shuts off one type of consumption entirely and leaves the other unchanged. The facts that consumption of the output that can still be produced remains at its usual level and that utility is additively separable imply that marginal utility in a pandemic is the same as in normal times. As a result, individuals don't require a premium to provide insurance against a pandemic state.

***The case of deficit-financed transfers.*** In practice, transfers in the pandemic have been financed not by current taxes but by borrowing, which will presumably be offset by higher taxes (or lower government spending) in the future. Thus, suppose that there are two periods; that a pandemic can only occur in the first period; that if there's a pandemic, the government finances any social insurance payments by taxes on all individuals in the second period; and that all individuals can save and borrow at the same interest rate as the government. Finally, assume an individual's utility is $[U(C_A^1) + U(C_B^1) - V(L^1)] + \beta[U(C_A^2) + U(C_B^2) - V(L^2)], 0 < \beta < 1$, where superscripts denote time periods.

In this situation, the complete-markets outcome is very similar to that of the baseline case: in the event of a pandemic (which can only occur in the first period by assumption), each individual's consumption of the output of the sector that stays open is ½, with no impact on allocations in the second period. For the government to bring about this outcome through first-period transfers and second-period taxes, it simply makes transfers to unemployed workers to fully offset their lost income in the first period, and then taxes everyone half their usual income in the second. With this policy, everyone saves half their after-transfer income in the first period and then uses their savings to pay the higher second-period taxes, and they all have the same path of consumption as under complete markets.

Notice that in this case, as in the baseline one, the higher consumption of the unemployed relative to what they would have in the absence of insurance or government intervention comes from lower consumption of the employed. In terms of income flows, however, the accounting is more complicated. The unemployed get transfers in excess of their consumption in the period of the pandemic, but they pay some of that back through second-period taxes; and the employed neither get transfers nor pay taxes in the period of the pandemic, but pay taxes in the second period.

Section E on the possibility of an aggregate demand shortfall considers an extension of the model to multiple periods that has more interesting implications.

***Introducing partial labor mobility*.** Allowing for some ex post labor mobility has relatively little effect in the baseline case. However, because it has more significant implications when the government doesn't have full information about who should optimally be working in a pandemic, we describe its impact in the baseline case.

To allow for some labor mobility, we assume that when a pandemic hits, each individual in the sector that shuts down draws a cost of switching sectors. The distribution of the cost is assumed to be continuous, and to be wide enough that some but not all individuals switch out of the sector that shuts. For simplicity, the cost is a direct utility cost (for example, stress or inconvenience), rather than a cost in terms of output (for example, having to hire movers to move to a new location).

In the complete-markets outcome with these assumptions, individuals only purchase insurance against the possibility that a pandemic hits their sector *and* that their moving cost is above some threshold. If a pandemic hits their sector and they drew a low moving cost, they switch sectors and receive no insurance payments. Thus with complete markets, partial labor mobility causes output in the sector that stays open to rise in a pandemic. However, consumption of the sector's output is fully equalized across individuals, as before.[12] All of this can be replicated

[12] If the costs of moving take the form of output rather than utility, individuals purchase insurance to cover any moving costs they incur, and what is equalized across individuals is consumption net of any moving costs.

through targeted government transfers (targeted in this case to workers in the sector that was shut with sufficiently high moving costs, rather than to all workers in that sector).

***Some variations that have little effect on the main messages.*** Many of the assumptions of the baseline case serve only to simplify the analysis and don't affect the main messages:

- It's easy to introduce a fourth state of the world where there's a pandemic in both sectors and so neither produces output. Under any allocation mechanism, the outcome in that state is trivial: no one works, there's no output, and everyone has zero consumption.
- It's not necessary to assume there are only two sectors and a pandemic causes half the economy to shut down. One way to relax this assumption is to assume there's a large number $N$ of sectors; that (as before) they're symmetric; and that a pandemic shuts fraction $f$ of the sectors, with the sectors that are shut chosen at random. Under these assumptions, all the previous results continue to hold, with the obvious changes: in the no-pandemic state, each individual's consumption of the output of each sector is $1/N$ (rather than ½); in a pandemic, the optimal social insurance policy is to tax each individual who remains employed fraction $f$ of their income (rather than half their income) and transfer the proceeds to the unemployed; and so on.
- Relaxing the assumption that the two sectors have the same probability of being shut by a pandemic complicates the equations but doesn't change anything of importance. For example, suppose only one sector faces a risk of being shut. Then individuals who work in that sector always have lower consumption than those who work in the other sector (to counterbalance the fact that they don't always have the disutility of working); and fewer individuals are in the sector that may shut than in the one that always stays open (because the sector that's at risk is on average less productive). But the key results of the baseline case carry over: if the complete-markets outcome is achieved through social insurance, that insurance takes the form of targeted transfers to the workers who lose

their jobs in a pandemic, with no "stimulus" and with transfers that don't fully replace the income that unemployed workers lose in a pandemic . In addition, as the probability of a pandemic approaches zero, outcomes converge to those of the baseline case.

- It's straightforward to relax the assumption that utility is additively separable. Suppose an individual's utility, rather than being $U(C_A) + U(C_B) - V(L)$, is $U(C_A, C_B) - V(L)$. Letting subscripts denote partial derivatives, we assume $U_1(\bullet) > 0$, $U_2(\bullet) > 0$, $U_{11}(\bullet) < 0$, and $U_{22}(\bullet) < 0$; we make no assumption about the sign of $U_{12}(\bullet)$ (though we assume $U_{11}(C_A, C_B)U_{22}(C_A, C_B) - [U_{12}(C_A, C_B)]^2 > 0$ for all non-negative $C_A$ and $C_B$). Finally, to maintain the symmetry of the sectors, we assume that $U(C_A, C_B) = U(C_B, C_A)$ for all non-negative $C_A$ and $C_B$. With this generalization, everything we say in the text continues to hold. The only substantive change is to the result described above about the actuarially fair pricing of insurance under complete markets. For example, suppose $U_{12}(\bullet) < 0$, which implies that when a pandemic shuts one sector, the marginal utility of consumption of the output of the sector that stays open is higher than in normal times. In this case, equilibrium marginal utility is higher than normal in a pandemic, and so there's a premium on the provision of insurance against a pandemic relative to the actuarially fair price. When $U_{12}(\bullet) > 0$—which seems less plausible—the opposite holds.

In the various extensions and variations on the model considered below, one can make analogous changes to the ones described here without changing the main messages.

### B. Incentives and Fairness

***The case considered in the text*.** In the text, we describe several reasons that it might not be feasible or desirable for the allocation in a pandemic to make individuals who continue to work worse off than those who stop working. Here, for concreteness we model these considerations by making the somewhat artificial assumption that everyone in the sector that remains open should work in the pandemic, but that which sector an individual is in isn't observable. With this assumption, for the individuals in the sector that's not affected by the pandemic to continue to work, they must not prefer claiming they're in the other sector to working. This condition is

(A1) $$U(C^E) - V(1) \geq U(C^U),$$

where $C^E$ and $C^U$ are the consumptions of employed and unemployed workers in the pandemic (of the output that's still being produced). Since $V(1) > 0$, this requires $C^E > C^U$.

One can show that the optimal allocation involves satisfying the incentive constraint (A1) with equality. Thus,

(A2) $$U(C^E) - V(1) = U(C^U).$$

The allocation must also satisfy the economy's resource constraint. Since total output in the sector that remains open is ½ and the number of individuals in each sector is ½, this condition is

(A3) $$C^E + C^U = 1.$$

Equation (A3) implies that total consumption is the same as in the baseline, and equation (A2) implies that the employed now consume more than the unemployed. It follows that the consumption of the employed is more than without the constraint, and the consumption of the unemployed is less. Introducing fairness or incentive considerations therefore reduces the optimal level of social insurance. As in the baseline case, the optimal policy can be implemented with targeted transfers to the unemployed and without any type of government stimulus: the government taxes employed workers amount $C^U$ and transfers the proceeds to the unemployed.

***Introducing partial labor mobility*.** Our baseline case assumes individuals are completely immobile between the two sectors in a pandemic. In reality, however, there were large flows of workers from industries that were largely shut by the pandemic to ones that weren't. And—in contrast to what we assume in our analysis of labor mobility in Section A—realistically the government has little information about which workers can move relatively easily.

To analyze these issues, we again model labor mobility by assuming that each individual in the sector that shuts draws a utility cost of switching sectors when a pandemic hits, and that the distribution of the cost is continuous and is such that some but not all workers move. Now, however, we assume the government only knows the distribution of the cost, and not each worker's draw. Thus the government has conflicting objectives. On the one hand, it wants to insure workers for whom switching would be very costly against large falls in their consumption. On the other hand, it wants to provide incentives for individuals with low switching costs to move.

One can show that this tradeoff leads the government to adopt policies with three consequences. First, and least interestingly, partial labor mobility mitigates the fall in output in a pandemic, as in the case where the government has complete information. This follows straightforwardly from the fact that some individuals switch to the sector that stays open. Second, workers who started in the sector that doesn't shut (and so remain employed there) obtain higher consumption than individuals who start in the sector that shuts and don't move. That is, social insurance doesn't fully insure the unemployed, as in the case with incentive or fairness constraints but no labor mobility. Third, and perhaps most interestingly, the optimal policy provides individuals who switch sectors with higher consumption than workers who are in the open sector throughout. That is, there's in effect a "moving bonus" to workers who switch sectors in a pandemic. The reason is that higher consumption of switchers, but not higher consumption of workers who stay in the open sector, increases the incentive to switch sectors.[13]

[13] This discussion presumes the government can do this. If it can't treat employed workers differently depending on which sector they were in originally (because of either limited information or fairness considerations), labor mobility still mitigates the fall in output, and the employed still consume more than the unemployed. This case is broadly similar to our baseline case of incentives and fairness, where allocations are characterized by just the two variables $C^E$ and $C^U$.

## C. High-Risk Essential Workers and Hazard Pay

***Assumptions***. As described in the text, we introduce high-risk essential workers by assuming there are three sectors rather than two, and that in the event of a pandemic both the utility of consuming one sector's output and the disutility of working in that sector rise. Let $U^H(C)$ be the utility of consuming the output of the high-risk essential sector in a pandemic, and $V^H(L)$ be the disutility of working in that sector in a pandemic. We assume $U^{H\prime}(C) > U'(C)$ for all $C \geq 0$, and $V^H(1) > V(1)$; we normalize $U^H(0)$ and $V^H(0)$ to 0. ($U(\bullet)$ continues to denote the utility from consuming the output of a sector that's not high-risk, and $V(\bullet)$ the disutility of working in such a sector.) As in our baseline case, ex ante the sectors are identical: each is equally likely to be the one that's shut by a pandemic, and equally likely to be the one that's essential in a pandemic. Individuals are again mobile ex ante but immobile ex post.

***High-risk essential workers in the baseline case***. If there aren't incentive or fairness considerations, adding an essential hazardous sector has little effect on the complete-markets outcome, and hence correspondingly little effect on optimal social insurance. As in the baseline case without a high-risk essential sector, the symmetry of the sectors causes individuals to be allocated equally among them. In the event of a pandemic, there's full employment in the two sectors that continue to operate (recall that we assume preferences are such that in equilibrium individuals always work unless prevented by a pandemic); and it's again straightforward to check that all individuals have equal consumption of the outputs of those sectors (for the same reasons as in the baseline case). As before, the complete-markets allocation can be replicated by the straightforward tax-and-transfer scheme of taxing individuals who remain employed what they would have spent on the output of the sector that is shut and giving the proceeds to the individuals who had been employed in that sector.

***Adding incentive or fairness considerations***. If allocations that make employed workers worse off than the unemployed aren't feasible or are viewed as undesirable, the implications of adding an essential high-risk sector are different. Denote the three sectors in a pandemic by $S$ ("shut down"), $H$ ("high-risk essential"), and $L$ ("lower-risk"), and the consumption of an individual in sector $j$ of the output of sector $i$ by $C_i^j$.). For workers in the high-risk sector, the condition that the employed must be at least as well off as the unemployed is

$$U^H(C_H^H) + U(C_L^H) - V^H(1) \geq U^H(C_H^S) + U(C_L^S). \tag{A4}$$

And for those in the lower-risk sector, it is

$$U^H(C_H^L) + U(C_L^L) - V(1) \geq U^H(C_H^S) + U(C_L^S). \tag{A5}$$

As in the case without a high-risk essential sector, the optimal allocation satisfies the constraints with equality. Since $V^H(1) > V(1) > 0$, it follows that individuals in the high-risk sector have higher consumption than individuals in the lower-risk sector, who in turn have higher consumption than individuals in the sector that shuts down.

As discussed in the text, the implication that individuals in the high-risk essential sector consume more than individuals in the lower-risk sector can be interpreted as hazard pay. (Note, however, that in the efficient allocation in the absence of incentive or fairness considerations, workers in both the high-risk sector and the open sector that's lower risk are taxed. As a result, hazard "pay" for individuals in the high-risk essential sector could take the form of reduced taxes rather than actual transfers.)

***Does the government need to provide the hazard pay*?** Suppose there aren't insurance markets or markets in Arrow-Debreu commodities, that prices and wages are determined after the state of the world is realized, and that there's potentially social insurance. In a pandemic, the outputs of the two sectors that remain open are equal, but the marginal utility of the output of the high-risk essential sector is greater than that of the output of the lower-risk sector. Thus if prices and wages are flexible, the price of the high-risk sector's output must be greater than that of the lower-risk sector's. This implies that the marginal revenue product of a worker in the high-risk essential sector is greater than that of a worker in the lower-risk sector, and hence that the wage in the high-risk sector is higher. Thus an element of hazard pay (in the sense of higher wages for high-risk essential workers) would arise endogenously.

There's no reason, however, for the rise in the wage in the high-risk essential sector to make individuals in the two sectors that are open exactly equally well off, which is what's required under the optimal social insurance policy. At a general level, since the economy involves taxes and transfers and features incentive or fairness constraints, there's no presumption that it will yield the efficient outcome. Specifically, if the difference in marginal utility between the outputs of the high-risk and lower-risk sectors is small and the gap in the disutility of working is large, the rise in the wage isn't enough to make high-risk workers as well off as lower-risk workers; and if the opposite pattern holds, the rise in the wage makes high-risk workers strictly better off than lower-risk workers. Thus government intervention is almost certainly necessary to ensure that high-risk essential workers receive the optimal amount of hazard pay. Finally, in actuality prices and wages appear to have been quite sticky in the pandemic, suggesting that in practice the amount of endogenously generated hazard pay was minimal.

## D. Heterogeneous Incomes and Self-Insurance

***Heterogeneous incomes in the static case.*** The most natural way to introduce income heterogeneity is to assume workers differ in their productivities. Concretely, consider the baseline case with the change that if individual $j$ works in a sector, they produce $\alpha_j > 0$ units of the sector's output (rather than producing 1 unit), with $\alpha$ varying across individuals. It's straightforward to show that in this case, each individual's consumption under complete markets scales with their $\alpha$. That is, if there's no pandemic, individual $j$ consumes $\alpha_j/2$ units of the output of each sector, and if there's a pandemic, they consume $\alpha_j/2$ units of the output of the sector that remains open. (The equilibrium also features the sum of the $\alpha$'s of the workers in each sector being equal, rather than the number of workers in each sector being equal.) Thus, social insurance that replicates the complete-markets outcome features payments to the unemployed proportional to their non-pandemic earnings.

The proportionality result continues to hold when there are incentive or fairness considerations. To see this, recall that in the absence of such considerations, optimal social insurance makes the ratio of any two individuals' marginal utilities constant across states of the world. Although that property no longer holds for all pairs of individuals when there incentive or fairness constraints, it still holds for any pair of individuals who are in the same sector. If we assume that utility takes the usual constant relative risk aversion form, $U(C) = C^{1-\theta}/(1-\theta), \theta > 0$, this result implies that the ratio of the consumptions of any two individuals in the same sector is constant across states. And this in turn implies that the consumption that optimal social insurance provides to the unemployed in a pandemic is proportional to their non-pandemic earnings.

***Introducing costs of social insurance and possibilities for self-insurance*.** Although the proportionality result holds both with and without incentive and fairness considerations, for simplicity we focus on the case without them in the remainder of this section. As discussed in the text, the result that optimal social insurance is strongly increasing in individuals' non-pandemic earnings is unlikely to hold if social insurance is costly and there are realistic possibilities for self-insurance. To investigate the implications of these factors, it's not necessary to specify and solve a general equilibrium model with less than actuarially fair insurance and some self-insurance through saving and borrowing. A partial equilibrium approach is sufficient to see the main points.

For concreteness, consider an individual who is in sector $A$. Let $C_{NO}^{A}$ be their consumption of sector-$A$ output in the event of a sector-$B$ pandemic in the case of no insurance (other than self-insurance), and let $C_{NO}^{B}$ be their consumption of sector-$B$ output in the event of a sector-$A$ pandemic in the no-insurance case. In the background, we think of there being a dynamic economy where individuals have some ability to use saving and borrowing to smooth their consumption in the face of shocks. $C_{NO}^{A}$ and $C_{NO}^{B}$ reflect what comes out of that setting in the absence of any possibilities for insurance beyond self-insurance. Here, however, we take them as given, and make no assumptions about them other than $C_{NO}^{A} \geq C_{NO}^{B}$—that is, that in the no-insurance case, the individual has higher consumption if a pandemic shuts the other sector than if a pandemic shuts their own sector.

Now suppose costly insurance is introduced: the individual can give up some consumption in the event of a pandemic in the other sector in exchange for greater consumption in the event of a pandemic in their own sector (where in each case consumption is of the output of the sector that stays open). Specifically, then can choose a quantity $X \geq 0$ by which to increase their consumption in the event of a sector-$A$ pandemic, at the cost of reducing their consumption in the event of a sector-$B$ pandemic by $(1+\lambda)X$, where $\lambda \geq 0$. $\lambda$ reflects the various costs of providing the insurance—what's typically called the "loading factor" of the insurance.

To see the effects of introducing costly insurance, it's useful to start by relating the situation we're considering to our baseline model (where there are no prospects for self-insurance, and where insurance is frictionless if it's present). In that case, $C_{NO}^{A} = 1$ (in the absence of insurance, in the event of a sector-$B$ pandemic the individual spends their entire income on the output of sector $A$); $C_{NO}^{B} = 0$ (in the absence of insurance, in the event of a sector-$A$ pandemic the individual earns nothing and consumes nothing); and $\lambda = 0$ (the pricing of insurance is actuarially fair). With the individual able to trade off consumption in the two pandemic states one-for-one, they buy insurance until their marginal utilities in the two pandemic states are equal. Given our assumptions about utility, this implies they have equal consumption in the two states. Thus they choose $X = ½$.

Now return to the case we're interested in. The individual chooses the amount of insurance, $X$, to solve

$$\max_{X \geq 0} U(C_{NO}^{B} + X) + U(C_{NO}^{A} - [1+\lambda]X). \tag{A6}$$

The solution satisfies

$$\text{(A7)} \quad \begin{cases} X = 0 & \text{if } \frac{U'(C_{NO}^B)}{U'(C_{NO}^A)} \leq 1 + \lambda \\ \frac{U'(C_{NO}^B + X)}{U'(C_{NO}^A - [1+\lambda]X)} = 1 + \lambda & \text{if } \frac{U'(C_{NO}^B)}{U'(C_{NO}^A)} > 1 + \lambda. \end{cases}$$

That is, if the marginal utilities in the absence of insurance in the two states are sufficiently close—specifically, if the ratio of the marginal utility the individual would have if a pandemic shut their sector to the marginal utility they would have if a pandemic shut the other sector is less than $1 + \lambda$—the individual buys no insurance. Otherwise, they buy until the ratio of marginal utilities is $1 + \lambda$. In this case, assuming the loading factor is strictly positive, it follows that they stop buying at a point where their consumption in a pandemic that hits their own sector ($C_{NO}^B + X$) is strictly less than their consumption in a pandemic that hits the other sector ($C_{NO}^A - [1 + \lambda]X$).

***Implications for social insurance*.** When insurance takes the form of government-provided social insurance, $\lambda$ reflects three factors. First, there's the cost of administering the program, verifying eligibility, making payments, suffering some losses from error and fraud, and so on. Second, the program causes distortions in individuals' behavior (most notably, in unemployed individuals' efforts to find jobs). And third, satisfying the government's intertemporal budget constraint requires raising more revenue at some point, which involves distortion costs.

Most discussions of the costs of providing unemployment insurance focus on the second factor—the distortionary effects on individuals' incentives. And as we discuss in the text, during much of the pandemic, unemployment insurance replacement rates were very high. However, multiple studies, using both individual and state-level data, find no evidence of any substantial negative effect of the high replacement rates on employment during the pandemic through their incentive effects on job-finding efforts (Altonji et al., 2020; Dube, 2021; Marinescu, Skandalis, and Zhao, 2021; and Ganong et al., 2021b). One reason is that, because workers didn't expect the high benefits to last, it appears they often preferred employment to unemployment even when the replacement rate exceeded 100 percent. A second reason is that because the main constraint on overall employment during the pandemic was from labor demand rather than supply, a modest negative effect of the high replacement rates on individuals' job-seeking had little impact on overall hiring.

The fact that the distortionary effects of unemployment insurance in a pandemic appear small (coupled with the fact that the program does not appear to involve large administrative and related costs relative to the scale of its spending) suggests that the main costs of providing unemployment insurance in a pandemic are likely to come from the third factor—the distortions associated with raising the additional revenue at some point to offset the unemployment insurance payments. Thus, $\lambda$ is likely to be slightly higher than the marginal distortion costs of raising revenue, which are typically thought to be about 0.3 or 0.4.

If utility is constant relative risk aversion $(U(C) = C^{1-\theta}/(1-\theta), \theta > 0)$, equation (A7) becomes

$$\text{(A8)} \qquad \begin{cases} X = 0 & \text{if } \frac{C_{NO}^B}{C_{NO}^A} \geq \frac{1}{(1+\lambda)^{1/\theta}} \\ \frac{C_{NO}^B + X}{C_{NO}^A - (1+\lambda)X} = \frac{1}{(1+\lambda)^{1/\theta}} & \text{if } \frac{C_{NO}^B}{C_{NO}^A} < \frac{1}{(1+\lambda)^{1/\theta}}. \end{cases}$$

Gandelman and Hernández-Murillo (2015) report that "[p]robably the most widely accepted" estimates of $\theta$ "lie between 1 and 3." A recent study by Calvet et al. (2021), however, suggests an average value around 5. And Chetty and Szeidl (2007) show that because there are substantial costs to changing some types of consumption quickly, the effective value of $\theta$ for the short run (which is what's relevant to unemployment insurance) may be much larger (see their equation [7] and Figure IV).

If the costs of unemployment insurance are large and risk aversion is low, the gap between the consumption of the unemployed and the employed under optimal insurance may be large. For example, if $\lambda$ = 0.6 and $\theta$ = 2 and the solution is interior, the individual would buy insurance to the point where their consumption if a pandemic shuts their sector is 21 percent less than their consumption if a pandemic shuts the other sector (since $1/1.6^{1/2} \approx 0.79$). Equivalently, under optimal but costly social insurance with these parameter values, in a pandemic the consumption of an individual in the sector that's shut is 21 percent lower than that of a comparable individual in the sector that's open. And if the gap would be smaller than 21 percent in the absence of any insurance other than self-insurance, it's optimal not to provide the individual with any social insurance at all. On the other hand, low costs of insurance and high risk aversion imply a small gap in consumption. For example, if $\lambda = 0.4$ and $\theta = 6$ (and again, if the solution is interior), in a pandemic the consumption of an individual in the sector that's shut is only 5 percent lower than that of a comparable individual in the sector that's open. Since the assumptions needed to obtain the 21 percent and 5 percent figures are somewhat extreme, this analysis suggests that 10 to 15 percent is a reasonable figure for the gap between the consumption of the unemployed and the employed under optimal but costly social insurance.

If the ability to self-insure is increasing in income, these results have two important implications for the progressivity of optimal social insurance. First, they tend to make optimal insurance payments rise less than proportionately with income. Second, they imply that if individuals with sufficiently high incomes have enough ability to self-insure that in the absence of social insurance, the ratio of their consumption if a pandemic shuts their sector to their consumption if a pandemic shuts the other sector is greater than the critical threshold, they shouldn't be provided with any social insurance.

Although our partial equilibrium analysis is enough to show these implications, a limitation of this approach is that it doesn't provide exact expressions for the amount of insurance individuals would choose to buy if the loading factor were $\lambda$—and thus, in our framework, for optimal social insurance payments. If an individual's self-insurance is sufficiently large that the consumption ratio we've been discussing is above the critical threshold (that is, if $C_{NO}^{B}/C_{NO}^{A} \geq 1/(1+\lambda)^{1/\theta}$), the optimal social insurance payment to the individual if a pandemic shuts their sector is indeed zero. But suppose that condition fails. In that situation, the individual won't necessarily devote social insurance payments solely to current consumption: because it's more efficient for the government to provide insurance than for individuals to self-insure, social insurance may in effect crowd out self-insurance. As a result, knowing only what the individual's consumption would be without social insurance isn't enough to determine the optimal insurance payment—the optimal payment also depends on the extent of crowding out. If self-insurance comes from borrowing and that borrowing is associated with a cost or wedge analogous to the loading factor of the social insurance, then the extent of the crowding out depends on the specifics of the wedge. For example, if the individual can borrow freely up to some limit but can't borrow more than that at any interest rate (so the wedge is zero up to the limit and then becomes infinite), there may be no crowding out. But if the individual faces a constant cost of borrowing that's larger than $\lambda$, there may be substantial crowding out.

## E. The Possibility of an Aggregate Demand Shortfall and a Need for Stimulus

As described in the text, our analysis of the aggregate demand effects of a pandemic is closely related to the analysis of Guerrieri et al. (2020). Specifically, consider a multi-period version of our baseline case where each period is described by the static version. Individual $j$'s lifetime utility is given by the natural extension of the utility function of the static model:

$$\mathcal{U}_j = \sum_t \beta^t \left[U\left(C_{jt}^A\right) + U\left(C_{jt}^B\right) - V(L_{jt})\right], \quad 0 < \beta < 1, \tag{A9}$$

where $t$'s denote time periods. Since we assume there's no capital in the economy (recall that each period is described by the static version of the model, with labor as the only input into production), net wealth is zero. However, we allow for the possibility of saving and borrowing at the individual level (that is, for trades among individuals of claims on future output). We assume all individuals are able to save, but some may be unable to borrow. Aside from whether they're able to borrow (and, as usual, which sector they're in), individuals are identical. Finally, for simplicity we assume that each individual starts with zero wealth, that a pandemic is only possible in the first period, and, again, that utility is constant relative risk aversion, $U(C) = C^{1-\theta}/(1-\theta), \theta > 0$.

In the absence of a pandemic, consumption is constant. The intertemporal Euler equation of individuals who can both borrow and save implies that the real interest rate, $r$, must satisfy $\beta(1+r)=1$, or

$$r=\frac{1-\beta}{\beta}. \tag{A10}$$

At that interest rate, individuals who cannot borrow don't want to, and so the borrowing constraint doesn't bind.

Now suppose there's a pandemic in period 1; for concreteness, assume it's in sector $A$. As we've seen, in the optimal allocation in our baseline case, each individual's consumption of the output of sector $B$ remains at its steady state level. Thus in this case, all individuals continue to satisfy their intertemporal Euler equation at the steady state interest rate given by (A10). That is, with optimal social insurance, aggregate demand falls by exactly as much as aggregate supply.

But suppose that because of incentive or fairness considerations or costly insurance, the optimal allocation in a pandemic involves greater consumption by the individuals who stay employed than by those who become unemployed. Since the efficient level of output in sector $B$ is unaffected, employed workers' consumption of the output of sector $B$ is above its steady state level, and unemployed workers' is below. At the steady state real interest rate, employed workers therefore want to save and unemployed workers want to borrow. In the absence of any borrowing constraints, there's no reason to expect an imbalance between the two in one direction rather than the other at that interest rate, and in fact the assumption of constant relative risk aversion utility implies they exactly balance.[14] But if some unemployed individuals cannot borrow, borrowing at a given interest rate is lower. As a result, at the optimal level of output and the steady state real interest rate, saving now exceeds borrowing. Thus for the economy to attain that output, there

[14] To see this, note that individual $j$'s Euler equation relating their consumption of the output of sector $B$ in periods 1 and 2 is $C_{j2}^B/C_{j1}^B=[\beta(1+r)]^{1/\theta}$, which implies that total consumptions in the two periods satisfy $C_2^B=[\beta(1+r)]^{1/\theta}C_1^B$. Since the efficient levels of sector-$B$ output in the two periods are the same and the steady state interest rate satisfies $\beta(1+r)=1$, the condition $C_2^B=[\beta(1+r)]^{1/\theta}C_1^B$ holds at the efficient level of output and the steady state interest rate.

must be some source of additional aggregate demand, which could come from a reduction in the real interest rate or from fiscal policy.[15]

### F. Debt and Payment Forgiveness

***Assumptions*.** To address the possibility of forgiveness of payment obligations or outstanding debts, we return to our baseline model, but we suppose individual $j$ has to make a fixed payment of $F_j$ that doesn't generate any current consumption. As in the baseline case, individuals are otherwise identical (other than, as usual, their choice of sector). Thus if there were no possibility of a pandemic, individual $j$'s total consumption of all outputs would be $1 - F_j$ rather than 1. We take the sizes of the payments that individuals make as given, although realistically they would reflect past shocks and decisions that leave individuals in different financial situations.

We also make several smaller changes to the model. First, to avoid having to introduce any additional actors into the economy, we assume the sum of the $F_j$'s across individuals is zero—that is, some individuals receive fixed payments rather than making them. Second, to get results that are particularly easy to interpret, we again assume utility is constant relative risk aversion. Third, to be able to gauge the implied scale of any forgiveness, we assume that rather than two sectors there are many, and that a pandemic shuts down fraction $f$ of them chosen at random, where $0 < f < 1$. With $N$ sectors rather than two, individual's $j$'s utility is $\sum_{i=1}^{N} U(C_{ji}) - V(L_j)$, where $C_{ji}$ is their consumption of the output of sector $i$. In equilibrium, $1/N$ individuals are in each sector. Finally, we start by assuming that the probability of a pandemic is close to zero.

***Qualitative implications*.** In our baseline case without the fixed payment obligations (and with no income heterogeneity), optimal social insurance causes everyone's after-tax-and-

[15] Guerrieri et al. (2020) consider these issues in much more generality. One result that's particularly relevant to our analysis is that the fact that the fall in aggregate demand exactly equals the fall in aggregate supply in the absence of borrowing constraints depends on our assumption that utility is additively separable both over sectors and over time, which implies that the *intra*temporal elasticity of substitution between the outputs of the two sectors in a given period equals the *inter*temporal elasticity of substitution between a given sector's output in two periods. With a more general utility function, if the intratemporal elasticity is less than the intertemporal elasticity, aggregate demand falls more than aggregate supply even without borrowing constraints.

transfer income to fall by the same amount in a pandemic, and hence everyone's consumption to fall by the same amount. In the presence of fixed payment obligations, if a pandemic reduces output from 1 to $1 - f$, a policy that causes everyone's consumption to fall by the same amount causes individual $j$'s consumption to fall from $1 - F_j$ to $1 - F_j - f$.

This policy doesn't replicate the complete-markets outcome, however. To see this, note that in the complete-markets allocation, for any two individuals the ratio of their marginal utilities is the same across states. With constant relative risk aversion utility, this requires that the ratio of their consumptions is the same across states. Thus optimal social insurance causes each individual's consumption to fall by the same *proportion* in a pandemic rather than by the same *amount*. When a pandemic causes aggregate consumption to fall by fraction $f$, the complete-markets allocation therefore has individual $j$'s consumption fall from $1 - F_j$ to $(1 - f)(1 - F_j)$, or $1 - f - (1 - f)F_j$. This expression differs from $1 - f - F_j$ (what their consumption would be if each individual's consumption fell by amount $f$) by $fF_j$. Thus the complete-markets allocation—or the corresponding allocation achieved through social insurance—reduces individuals' fixed payment obligations by the same proportion that aggregate consumption falls. And in a multi-period version of the model where the pandemic is temporary, with complete markets individual $j$'s consumption reverts to $1 - F_j$ after the end of the pandemic. That is, the model points to partial payment forgiveness during the pandemic, not to permanent forgiveness of the underlying obligation that gives rise to the payments.

Relaxing the assumption that the probability of a pandemic is close to zero has little impact on these results. The place where this assumption enters the analysis is in the statement that if individual $j$ has a fixed payment obligation of $F_j$, their non-pandemic consumption is $1 - F_j$. This claim implicitly leaves out the individual's spending on insurance against a pandemic; this is appropriate when the chances of a pandemic are negligible, but not otherwise. When the probability of a pandemic isn't negligible, neither is spending on insurance against a pandemic.

As a result, in the general case the fact that individuals with larger $F_j$'s buy more insurance causes individual *j*'s non-pandemic consumption to take the form $1-\gamma(p)F_j$, where $p$ is the probability of a pandemic and where $\gamma(0)=1, \gamma'(\bullet)>0$. In this case, individual *j*'s consumption in a pandemic is $(1-f)(1-\gamma(p)F_j)$.

***Quantitative implications*.** At its low in April 2020, real personal consumption expenditures (PCE) in the U.S. were 18 percent below their peak level. And with the exception of April and May 2020, they were never more than 7 percent below their peak.[16] Thus a baseline model of optimal social insurance would have called for forgiving roughly 10 percent of fixed payment obligations during the pandemic. And it wouldn't have suggested a reason for forgiving outstanding debt balances. Thus someone with, for example, student debt would have had their payments reduced by roughly 10 percent during the pandemic and then returned to their pre-pandemic level when the pandemic ended.

***The size of the welfare gains*.** It's natural to wonder whether the welfare gains from the optimal forgiveness policy are large. Under no forgiveness but otherwise optimal social insurance (so each individual's consumption falls by amount *f* in a pandemic), individual *j*'s expected utility is

$$\text{(A11)} \qquad \mathcal{U}_{NO,j} = (1-p)\left[NU\left(\frac{1-F_j}{N}\right) - V(1)\right] + p\left[(1-f)NU\left(\frac{1-f-F_j}{(1-f)N}\right) - (1-f)V(1)\right].$$

Under the complete-markets allocation, which includes reductions in fixed payment obligations in a pandemic, it's

$$\text{(A12)} \qquad \mathcal{U}_{FORG,j} = (1-p)\left[NU\left(\frac{1-\gamma(p)F_j}{N}\right) - V(1)\right] + p\left[(1-f)NU\left(\frac{(1-f)(1-\gamma(p)F_j)}{(1-f)N}\right) - (1-f)V(1)\right].$$

As discussed above, $\gamma(p)$ reflects the fact that with complete markets, individuals with higher fixed payments buy more insurance against a pandemic. Recall that in the complete-markets

[16] Looking at real PCE on nondurables and services, which are probably more relevant to current utility than real total PCE, yields a similar picture. The decline in real GDP (which is the same as the decline in real consumption in the model, since all output is consumed) was smaller than the decline in real PCE.

allocation, each individual's marginal utility is no different in a pandemic than in normal times, and so the pricing of pandemic insurance is actuarially fair. It follows that the expected shortfall of individual $j$'s spending on consumption from their income must equal their payment obligation, $F_j$. That is, the individual can use insurance to shift payments between the non-pandemic and pandemic states, but can't change their expected value. In the non-pandemic state, the individual's fixed payment is $\gamma(p)F_j$, and in the pandemic state, it's $\gamma(p)(1-f)F_j$. Thus their expected payment is $\gamma(p)[(1-p)+(1-f)p]F_j$, or $\gamma(p)(1-fp)F_j$. Since this must equal $F_j$, it follows that

(A13) $$\gamma(p)=\frac{1}{1-fp}.$$

The change in individual $j$'s expected utility due to the optimal forgiveness policy is $\Delta\mathcal{W}_j \equiv \mathcal{U}_{FORG,j}-\mathcal{U}_{NO,j}$. To convert this from units of utility to units of consumption, we can divide it by the individual's expected marginal utility of consumption in the absence of forgiveness, which we denote $\mathcal{U}'_{NO,j}$. A natural measure of the overall welfare benefit of the optimal forgiveness policy is the average of this measure across individuals.[17] Thus our measure of the benefit is

(A14) $$\Delta\mathcal{W} \equiv \int_{j=0}^{1}[(\mathcal{U}_{FORG,j}-\mathcal{U}_{NO,j})/\,\mathcal{U}'_{NO,j}]dj.$$

We take a second-order Taylor approximation of (A14) around the point where all the $F_j$'s are zero. The result is:

(A15) $$\Delta\mathcal{W} \cong \frac{1}{2}\theta\frac{f^2p(1-p)}{(1-f)(1-fp)}Var(F_j),$$

where $\theta$ is the coefficient of relative risk aversion. Dividing this expression by the probability of a

[17] An alternative is to compute the average gain in expected utility across individuals, and then normalize by average expected marginal utility across individuals. For the second-order approximation that we consider, this approach yields the same expression for the welfare benefit of optimal forgiveness as what we find below. A disadvantage of using the measure based on average expected utility, however, is that it implies that the optimal policy is simply to forgive all fixed payment obligations entirely, even in the non-pandemic state. Thus using this welfare measure raises issues unrelated to the appropriate response to a pandemic.

pandemic gives the welfare gain per pandemic (that is, loosely speaking, the welfare gain if a pandemic occurs):

$$\Delta\mathcal{W}/p \cong \frac{1}{2}\theta\frac{f^2(1-p)}{(1-f)(1-fp)}Var(F_j). \quad \text{(A16)}$$

As we describe above, $f$ = 0.1 is a generous estimate of the fraction of the economy that was shut in most of the pandemic. Similarly, 20 percent of normal consumption is probably a generous estimate of the standard deviation of fixed payment obligations across individuals. If, in addition, we assume $p$ = 0.01 (which has little impact on the results over a reasonable range) and $\theta$ = 4, we obtain $\Delta\mathcal{W}/p \cong 0.00088$. Since normal consumption is 1, this means that the welfare benefit, per pandemic, of the optimal forgiveness policy is roughly 0.09 percent of normal consumption, which is very small.

This discussion assumes there are no frictions or administrative costs to providing partial forgiveness. In practice, determining precisely what constitutes a fixed payment obligation, measuring each individual's obligations, and implementing a program of temporary partial forgiveness would be challenging. Given the small welfare gains from forgiveness under the optimal policy in the absence of administrative costs, with plausible costs the welfare gains might well be negative. In short, a social insurance perspective suggests at most relatively little forgiveness in a pandemic, and under realistic assumptions it might point to no forgiveness at all.

***Forbearance*.** An alternative to forgiveness is forbearance—allowing individuals to not make their fixed payment obligations during the pandemic, and instead make them, together with the accrued interest, when conditions return to normal. Thus forbearance effectively gives individuals with fixed payment obligations a way to borrow. (A third possibility, forbearance without having to pay interest, amounts to forbearance plus partial forgiveness.)

With fixed payment obligations and the optimal forgiveness policy, each individual's marginal utility in a pandemic is the same as in non-pandemic times. As a result, the marginal benefit (if any) to borrowing during a pandemic for individuals with large fixed payment obligations is no greater than in non-pandemic times. Thus there's no more reason to provide forbearance in a pandemic than in other times.

To see this concretely, consider a multi-period version of our baseline case where each period is the same as the static version, and individual $j$'s fixed payment obligation is $F_j$ every period (except in a pandemic, when it's $(1-f)F_j$). Then—along the lines of Section E—each individual's marginal utility under optimal policy is constant over time, no one wants to borrow or save, and so there are no benefits to forbearance.

As usual, however, the situation is different if incentive or fairness considerations cause the unemployed to have lower consumption than the employed in a pandemic. In a multi-period setting, if the unemployed face borrowing constraints, there might be scope for Pareto improvements by using forbearance to allow unemployed workers with fixed payment obligations to shift consumption from the future to the present. Thus there might be a case for forbearance in this situation. There are two important caveats, however. First, the fixed payment obligations have no direct role in these potential welfare benefits—there could be welfare benefits from allowing all unemployed individuals, regardless of the size of their fixed payment obligations, to in effect borrow against their future income. Second, the same logic that points to forbearance in a pandemic also points to policies that would allow the unemployed to borrow in normal times. Thus, a full analysis of forbearance would require addressing the issues of why the unemployed face borrowing constraints in private markets, and why there aren't currently large-scale policies of government-provided loans to the unemployed in normal times.

## G. Aid to State and Local Governments

A natural extension of our social insurance framework can help us understand a somewhat different issue than aid to individuals in a pandemic—the possibility of federal aid to state and local governments. To consider the treatment of state and local governments (which we refer to as "local governments" for simplicity), we again start from our baseline case. We introduce a local government sector by assuming individuals obtain utility from the outputs of the two sectors, $A$ and $B$, and from the output of local governments. We continue to assume that utility is additively separable. Thus an individual's utility is $U(C_A) + U(C_B) + H(G) - V(L)$, where $G$ is the output of the local government sector, where we assume $H'(\bullet) > 0$ and $H''(\bullet) < 0$, and where we normalize $H(0)$ to 0. We think of $G$ as a public good. Its production function is $G = L_G$, where $L_G$ is the number of workers in the local government sector. As before, we assume workers are fully mobile across sectors ex ante but fully immobile ex post, and we assume sectors $A$ and $B$ are equally likely to be shut by a pandemic. In addition, we allow for the possibility that a pandemic makes the output of local governments more valuable, so that $H(\bullet)$ is replaced by $H^p(\bullet)$ in a pandemic, with $H^{p\prime}(\bullet) \geq H'(\bullet)$ for all $G$ (and with $H^p(0)$ again normalized to 0). Finally, for simplicity we assume the probability of a pandemic is small.

Under these assumptions, the efficient outcome is for workers to be allocated across sectors so that the marginal utility of output of each sector if there isn't a pandemic is the same. This obviously requires equal numbers of workers in sectors $A$ and $B$. Since $L_A$, $L_B$, and $L_G$ must sum to 1, $L_A$ and $L_B$ must therefore each equal $(1 - L_G)/2$. It follows that the condition for the optimal number of workers in the government sector is[18]

$$H'(L_G) = U'\left(\frac{1-L_G}{2}\right). \tag{A17}$$

A pandemic has no effect on the utility cost of the marginal worker continuing to work in the government sector, and either raises or has no effect on the utility benefit. Thus the efficient allocation involves continued full employment in the local government sector.[19] Thus if

[18] This is the only place where the assumption that the probability of a pandemic is small enters this analysis. In the case of a general $p$, the condition for the optimal allocation of workers to the local government sector is more complicated than (A17), but approaches (A17) as $p$ approaches 0. This has no substantive implications.

[19] The reason we refer to this outcome as "efficient" rather than "complete-markets" is that because government output is a public good, obtaining it requires not just complete markets but also some appropriate political process. Since individuals are ex ante identical, the allocation we describe is preferred unanimously to any other allocation that doesn't provide higher expected utility to some individuals than

individuals contemplated the possibility of a pandemic, they would likely agree to have their local governments purchase pandemic insurance, or to have automatic tax increases on workers who stay employed in a pandemic.

In the absence of fully specified ex ante agreements about allocations, however, implementing the efficient outcome may not be feasible. In practice, local governments are subject to balanced budget requirements that must be satisfied conditional on the realized state of the world. And local governments don't have policies that make taxes rise automatically in a pandemic. Further, although raising taxes ex post in the face of a pandemic could implement the efficient allocation, an across-the-board tax increase wouldn't be a Pareto improvement conditional on the realized state. For example, local government workers who would remain employed even without the tax increase would be worse off. Thus under realistic political assumptions, taxes might not change immediately, in which case a pandemic would lead to a fall in tax revenues, and so through the balanced budget requirement to an inefficient fall in $G$. This points to the potential for a government that can borrow—in the U.S. case, the federal government—to step in and prevent the inefficient fall in local government output.

## H. <u>Incentives and Fairness with Partial Information</u>

The analysis of incentive and fairness considerations in the text and in Section B of this appendix considers the extreme case where the government has no information at all about who should optimally be working in a pandemic. An intermediate possibility is that it has some but less than full information. For example, the government may know that the optimal allocation involves a smaller fraction of older individuals than younger individuals working in a pandemic, but not know precisely which older and younger individuals should work. This section therefore provides a brief discussion of such cases.

---

others. Thus it seems reasonable to assume that if there's a political process for determining state-contingent outcomes ex ante, that allocation would be the outcome.

A straightforward way to allow for partial information is to extend the analysis in the first part of Section B to the case of multiple groups. The government knows which group each person is in and the fraction of each group that will work in a pandemic in the optimal allocation, but not who within each group should be working. In contrast to our analysis of partial labor mobility and paralleling what we do in the first part of Section B, we assume that those who should be working face no cost of doing so other than the usual disutility of work, and that working is prohibitively costly for those who shouldn't be working.

To get a sense of the implications of these assumptions, we consider a simple case with three types of individuals. The first type are all individuals who should work in a pandemic, and the other two types are all individuals who shouldn't be working in a pandemic. The government, however, cannot distinguish between the first two types, and so from its point of view individuals fall into two groups: one (the first two types) consists of a mix of individuals who should and shouldn't be working, and one (the third type) consists entirely of individuals who shouldn't be working. For concreteness, imagine that whether an individual should work in a pandemic depends only on whether they have health risks (and not what sector they were in originally), but that some health risks are observable (those of individuals of the third type) and some are not (those of individuals of the second type).

We denote the three types by "*E*" ("employed"), "*U*" ("unobservable risk"), and "*O*" ("observable risk"). Paralleling our assumptions in the baseline case, we assume that ex ante all individuals have the same chances as one another of being in each category. $C^i$ denotes the consumption of a representative individual of type $i$, and $f_i$ denotes the probability of being of type $i$. We assume the $f_i$'s (which must sum to 1) are all strictly positive.

These assumptions imply that everyone's expected utility in a pandemic ex ante (that is, before they know which type they are) is

$$\mathcal{U} = f_E[U(C^E) - V(1)] + f_U U(C^U) + f_O U(C^O). \tag{A17}$$

The optimal allocation involves choosing the $C$'s to maximize $\mathcal{U}$ subject to two constraints. The first is that, since the government cannot distinguish individuals of the first two types, for individuals of the first type working must be at least as attractive as not working and having the consumption of individuals of the second type. This condition is

$$U(C^E) - V(1) \geq U(C^U). \tag{A18}$$

As before, the optimal allocation satisfies the incentive constraint with equality. The second constraint is the economy's resource constraint:

$$f_E C^E + f_U C^U + f_O C^O = f_E. \tag{A19}$$

The main message that comes out of analyzing this case is that the optimal allocation has $C^E > C^O > C^U$. That is, individuals who the government knows for sure shouldn't be working (the "$O$" type) receive larger transfers than the individuals in the group the government is unsure about but who in fact shouldn't be working (the "$U$" type); but these transfers aren't enough to raise the consumption of the "$O$" type to that of the employed (the "$E$" type). This result is broadly similar to our finding in the second part of Section B that with some labor mobility and imperfect information, individuals who the government knows for sure should be working obtain consumption between that of those who are induced to switch sectors and those who are unemployed.

One can also show that the optimal $C^O$ doesn't necessarily equal $f_E$, the consumption that everyone would obtain in the baseline case where the government has full information about who should and shouldn't be working. It turns out that if the utility function, $U(\bullet)$, is logarithmic, the optimal $C^O$ is exactly $f_E$, but that in the general case it can be either more or less than $f_E$.

It's conceptually straightforward to extend this analysis to $N$ groups of potentially different sizes, with different fractions of individuals who should optimally be working in a pandemic.

**REFERENCES**


Altonji, Joseph, Zara Contractor, Lucas Finamor, Ryan Haygood, Ilse Lindenlaub, Costas Meghir, Cormac O'Dea et al. 2020. "Employment Effects of Unemployment Insurance Generosity during the Pandemic." Unpublished paper (July).

Baker, Scott R., R. A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. 2021. "Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments." Unpublished paper (February).

Billingsley, Sunnee, Maria Brandén, Siddartha Aradhya, Sven Drefahl, Gunnar Andersson, and Eleonora Mussino. 2020. "Deaths in the Frontline: Occupation-Specific COVID-19 Mortality Risk in Sweden." Stockholm Research Reports in Demography 2020:36.

Bitler, Marianne P., Hilary W. Hoynes, and Diane Whitmore Schanzenbach. 2020. "The Social Safety Net in the Wake of COVID-19." *Brookings Papers on Economic Activity* (Summer): 119–145.

Blau, Francine D., Josefine Koebe, and Pamela A. Meyerhofer. 2021. "Who Are the Essential and Frontline Workers?" *Business Economics* 56 (July): 168–178.

Cajner, Tomaz, Andrew Figura, Brendan M. Price, David Ratner, and Alison Weingarden. 2020. "Reconciling Unemployment Claims with Job Losses in the First Months of the COVID-19 Crisis." Federal Reserve Board, Finance and Economics Discussion Series 2020-055 (July).

Calvet, Laurent E., John Y. Campbell, Francisco Gomes, and Paolo Sodini. 2021. "The Cross-Section of Household Preferences." NBER Working Paper No. 28788 (May).

Centers for Disease Control and Prevention. 2021a. "Estimated Disease Burden of COVID-19." (April).

Centers for Disease Control and Prevention. 2021b. "Weekly Updates by Select Demographic and Geographic Characteristics." Data as of 4/28/2021.

Chen, Yea-Hung, Maria Glymour, Alicia Riley, John Balmes, Kate Duchowny, Robert Harrison, Ellicott Matthay, Kirsten Bibbins-Domingo. 2021. "Excess Mortality Associated with the COVID-19 Pandemic among Californians 18–65 Years of Age, by Occupational Sector and Occupation: March through November 2020." *PLOS ONE* 16 (June 4): e0252454.

Chetty, Raj, and Adam Szeidl. 2007. "Consumption Commitments and Risk Preferences." *Quarterly Journal of Economics* 122 (May): 831–877.

Chodorow-Reich, Gabriel, and John Coglianese. 2019. "Unemployment Insurance and Macroeconomic Stabilization." In *Recession Ready: Fiscal Policies to Stabilize the American Economy*. Heather Boushey, Ryan Nunn, and Jay Shambaugh, eds. (Washington, D.C.: Brookings Institution): 153–179.

Dube, Arindrajit. 2021. "Aggregate Employment Effects of Unemployment Benefits during Deep Downturns: Evidence from the Expiration of the Federal Pandemic Unemployment Compensation." NBER Working Paper No. 28470 (February).

Duncan, Greg J. 1977. "Earnings Functions and Nonpecuniary Benefits." *Journal of Human Resources* 11 (November): 462 –483.

Farrell, Diana, Peter Ganong, Fiona Greig, Max Liebeskind, Pascal Noel, and Joseph Vavra. 2020. "Consumption Effects of Unemployment Insurance during the COVID-19 Pandemic." Technical report. JPMorgan Chase Institute (July).

Gandelman, Néstor, and Rubén Hernández-Murillo. 2015. "Risk Aversion at the Country Level." Federal Reserve Bank of St. Louis *Review* 97 (First Quarter): 53–66.

Ganong, Peter, Fiona Greig, Max Liebeskind, Pascal Noel, Daniel M. Sullivan, and Joseph Vavra. 2021a. "Spending and Job Search Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data." Unpublished paper (February).

Ganong, Peter, Fiona Greig, Pascal Noel, Daniel M. Sullivan, and Joseph Vavra. 2021b. "Micro and Macro Disincentive Effects of Expanded Unemployment Benefits." Unpublished paper (July).

Ganong, Peter, and Pascal Noel. 2019. "Consumer Spending during Unemployment: Positive and Normative Implications." *American Economic Review* 109 (7) (July): 2383–2424.

Ganong, Peter, Pascal Noel, and Joseph Vavra. 2020. "US Unemployment Insurance Replacement Rates during the Pandemic." *Journal of Public Economics* 191 (November): 104273.

Greenhalgh, Trisha, Matthew Knight, Christin A'Court, Maria Buxton, and Laiba Husain. 2020. "Management of Post-Acute Covid-19 in Primary Care." *British Medical Journal* 370: m3026 (August 11).

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2020. "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" Unpublished paper (July).

Karger, Ezra, and Aastha Rajan. 2021. "Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Payments." Federal Reserve Bank of Chicago Working Paper 2020-15 (revised) (February).

Kearney, Audrey, and Cailey Muñana. 2020. "Taking Stock of Essential Workers." Kaiser Family Foundation Policy Watch (May 1).

Kinder, Molly, Laura Stateler, and Julia Du. 2020. "The COVID-19 Hazard Continues, but the Hazard Pay Does Not: Why America's Essential Workers Need a Raise." Brookings Metropolitan Policy Program Report (July).

Magnusson, Karin, Karin Nygård, Line Vold, and Kjetil Telle. 2021. "Occupational Risk of COVID-19 in the 1st vs 2nd Wave of Infection." Preprint, Norwegian Institute of Public Health (January 6).

Marinescu, Ioana, Daphne Skandalis, and Daniel Zhao. 2021. "The Impact of the Federal Pandemic Unemployment Compensation on Job Search and Vacancy Creation." NBER Working Paper No. 28567 (March).

McLernon, Lianna Matt. 2020. "COVID-Related Nursing Shortages Hit Hospitals Nationwide." *Center for Infectious Disease and Research Policy News* (November 30).

Milne, Alistair. 2020. "A Critical COVID-19 Economic Policy Tool: Retrospective Insurance." Unpublished paper (March).

Mutambudzi, Miriam, Claire Niedzwiedz, Ewan Beaton Macdonald, Alastair Leyland, Frances Mair, Jana Anderson, Carlos Celis-Morales et al. 2021. "Occupation and Risk of Severe COVID-19: Prospective Cohort Study of 120 075 UK Biobank Participants." *Occupational and Environmental Medicine, British Medical Journal* 78 (April 15): 307–314.

Nguyen, Daisy. 2020. "US Hospitals Facing Worrisome Shortages of Nurses, Doctors." APNews.com (December 22).

O'Driscoll, Megan, Gabriel Ribeiro Dos Santos, Lin Wang, Derek A. T. Cummings, Andrew S. Azman, Juliette Paireau, Arnaud Fontanet, Simon Cauchemez, and Henrik Salje. 2020. "Age-Specific Mortality and Immunity Patterns of SARS-CoV-2." *Nature* 590 (February 4): 140–145.

Pennsylvania Department of Community and Economic Development. 2020. "COVID-19 PA Hazard Pay Grant Program." Program Guidelines (July).

Rowe, Nick. 2020. "Relative Supply Shocks, Unobtainium, Walras' Law, and the Coronavirus." Worthwhile Canadian Initiative: A Mainly Canadian Economics Blog (March).

Saez, Emmanuel, and Gabriel Zucman. 2020. "Keeping Businesses Alive: The Government Will Pay." Unpublished paper (March).

Song, Hummy, Ryan M. McKenna, Angela T. Chen, Guy David, and Aaron Smith-McLallen. 2021. "The Impact of the Non-Essential Business Closure Policy on Covid-19 Infection Rates." *International Journal of Heath Economics and Management* (April 1).

U.S. Bureau of Labor Statistics. 2019. "May 2019 National Occupational Employment and Wage Estimates, United States." https://www.bls.gov/oes/2019/may/oes_nat.htm.

U.S. Census Bureau. 2020. "Working-Age Population Not Keeping Pace with Growth in Older Americans. https://www.census.gov/library/stories/2020/06/working-age-population-not-keeping-pace-with-growth-in-older-americans.html (June).

U.S. Congress. 2020. "H. R. 6800." 116th Congress, 2d Session (May).

U.S. Environmental Protection Agency. 2016. "Valuing Mortality Risk Reductions for Policy: A Meta-Analytic Approach." White paper.

U.S. Office of Personnel Management. (2021). "Fact Sheet: Computing Hourly Rates of Pay Using the 2,087-Hour Divisor. https://www.opm.gov/policy-data-oversight/pay-leave/pay-administration/fact-sheets/computing-hourly-rates-of-pay-using-the-2087-hour-divisor/.

Weinzierl, Matthew. 2014. "The Promise of Positive Optimal Taxation: Normative Diversity and a Role for Equal Sacrifice." *Journal of Public Economics* 118 (October): 128–142.

Woodford, Michael. 2020. "Effective Demand Failures and the Limits of Monetary Stabilization." NBER Working Paper No. 27768 (September).