Macroeconomic Implications of COVID-19:

Can Negative Supply Shocks Cause Demand Shortages?

Veronica Guerrieri

Chicago Booth

Guido Lorenzoni

Northwestern

Ludwig Straub

Harvard

Iván Werning

MIT

April 2, 2020

We present a theory of *Keynesian supply shocks*: supply shocks that trigger changes in aggregate demand larger than the shocks themselves. We argue that the economic shocks associated to the COVID-19 epidemic—shutdowns, layoffs, and firm exits—may have this feature. In one-sector economies supply shocks are never Keynesian. We show that this is a general result that extend to economies with incomplete markets and liquidity constrained consumers. In economies with multiple sectors Keynesian supply shocks are possible, under some conditions. A 50% shock that hits all sectors is not the same as a 100% shock that hits half the economy. Incomplete markets make the conditions for Keynesian supply shocks more likely to be met. Firm exit and job destruction can amplify the initial effect, aggravating the recession. We discuss the effects of various policies. Standard fiscal stimulus can be less effective than usual because the fact that some sectors are shut down mutes the Keynesian multiplier feedback. Monetary policy, as long as it is unimpeded by the zero lower bound, can have magnified effects, by preventing firm exits. Turning to optimal policy, closing down contact-intensive sectors and providing full insurance payments to affected workers can achieve the first-best allocation, despite the lower per-dollar potency of fiscal policy.

1 Introduction

Jean-Baptiste Say is famously misquoted for stating the Law "supply creates its own demand." In this paper, we introduce a concept that might be accurately portrayed as "supply creates its own *excess* demand". Namely, a negative supply shock can trigger a demand shortage that leads to a contraction in output and employment larger than the supply shock itself. We call supply shocks with these properties *Keynesian supply shocks*.

Temporary negative supply shocks, such as those caused by a pandemic, reduce output and employment. As dire as they may be, supply shock recessions are partly an efficient response, since output and employment should certainly fall. However, can a supply shock induce too sharp a fall in output and employment, going beyond the efficient response? Can it lead to a drop in output and employment for sectors that are not directly affected by shutdowns? Relatedly, could this process produce an anemic recovery or is a V-shaped recession assured?

These are the questions we seek to address in this paper. They are also the questions behind recent debates over monetary and fiscal policy responses to the COVID-19 epidemic and the ensuing economic fallout. We also examine the logic behind these class of stabilization measures.

A simple perspective on the effects of COVID-19, casts the issue as one of aggregate supply versus aggregate demand, whether the shock to one side is greater than the other. Some have expressed skepticism that any demand stimulus is warranted in response to what is essentially a supply shock, and argue that the economic response should be purely framed in terms of social insurance. Others have expressed the belief that the pandemic shock can cause output losses larger than efficient. For example, Gourinchas (2020) has argued for macro measures aimed at "flattening the recession curve." The debate illustrates that a discussion focused on demand versus supply opens up many possibilities, but leaves many questions unanswered. What forces would induce demand to contract more than supply?

The perspective we offer here is different and based on the notion that supply and demand forces are intertwined: demand is endogenous and affected by the supply shock and other features of the economy. Our analysis uncovers features of the economy that matter and the mechanisms by which forces acting on the supply side end up affecting the demand side as well. The basic intuition is simple: when workers lose their income, due to the shock, they reduce their spending, causing a contraction in demand. However, the question is whether this mechanism is strong enough to cause an overall shortfall in demand.

Figure 1: How negative supply shocks can lead to demand shortages
— Case with equal inter- and intra-temporal elasticities —
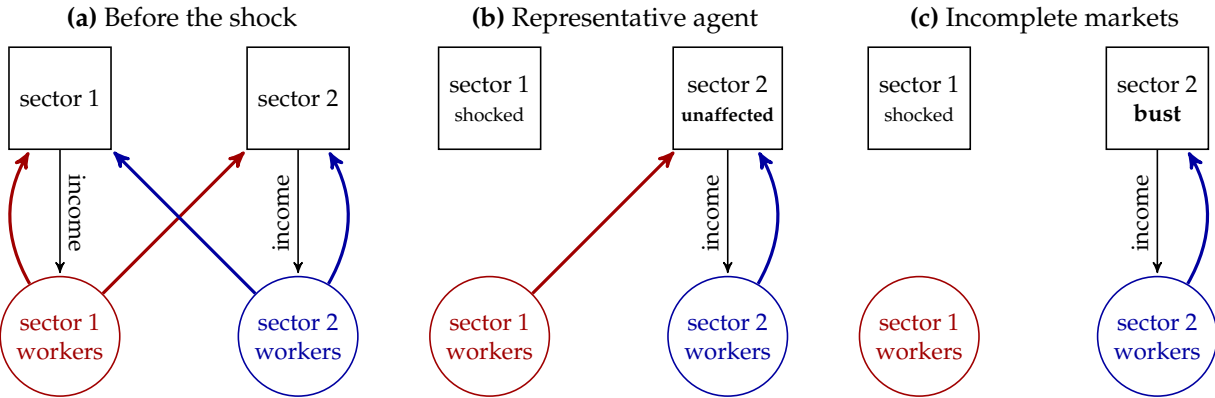
First, we show that in one-sector economies the answer is negative: the drop in supply dominates. The result is well-known in a representative agent economy. Less obviously, we show that it holds true in richer incomplete market models that allow for heterogeneous agents, uninsurable income risk and liquidity constraints, creating differences in marginal propensities to consume (MPC). In these models, a mechanism from income loss to lower demand is present, but although it makes the drop in aggregate demand larger than in the representative agent case, the drop is still smaller than the drop in output due to the supply shock. Intuitively, the MPCs of people losing their income may be large, but is bounded above at one, which implies that their drop in consumption is always a dampened version of their income losses.

We then turn to economies with multiple sectors. When shocks are concentrated in certain sectors, as they are during a shutdown in response to an epidemic, there is greater scope for total spending to contract. The fact that some goods are no longer available makes it less attractive to spend overall. An interpretation is that the shutdown increases the shadow price of the goods in the affected sectors, making total current consumption more expensive and thus discouraging it. On the other hand, the unavailability of some sectors's goods can shift spending towards the other sectors, through a substitution channel. Whether or not full employment is maintained in the sectors not directly affected by the shutodwn depends on the relative strength of these two effects.

We show that a contraction in employment in unaffected sectors is possible in a representative agent setting when the intertemporal elasticity of substitution is sufficiently

high and the elasticity of substitution across sectors is not too large.[1] An alternative intuition is that under these conditions the two goods are Hicks complements, so that lower marginal consumption of goods affected by the shutdown decreases the marginal utility from consuming unaffected goods.

We then turn to incomplete markets and show that the condition for a contraction in employment in unaffected sectors becomes less stringent. Intuitively, if workers in the affected sectors lose their jobs and income, their consumption drops significantly if they are credit constrained and have high MPCs. To make up for this, workers in the unaffected sectors would have to increase their consumption of the remaining goods sufficiently. This requires a higher degree of substitution across sectors. If goods are not too close substitutes, aggregate demand contracts more than supply and employment in the unaffected sectors falls.

Figure 1 illustrates this logic for two sectors, 1 and 2, where sector 1 gets shocked. In a representative agent setting, agents working in both sectors pool their income and spend it across sectors identically. Here, the difference between inter- and intra-temporal elasticities matters for whether sector 2 is affected by the shock in sector 1. Figure 1(b) shows the knife-edge case where both elasticities are equal and sector 2 is unaffected. Panel (c) then emphasizes that with incomplete markets, even this case causes sector 2 to go into a recession, as sector 1 workers cut back their spending on sector 2. Thus, Figure 1 illustrates how a supply shock in sector 1 can spill over into a demand shortage in sector 2, that is amplified by incomplete markets.

The fact that aggregate demand causes a recession above and beyond the reduction in supply might lead one to think that fiscal policy interventions are powerful in keeping aggregate demand up. We show that this is a false conclusion. First of all, the marginal propensity to consume may be low. Second, and more surprisingly, the standard Keynesian cross logic behind fiscal multipliers is not operational in the recession, there are no second round effects, so the multiplier for government spending is 1 and that for transfers is less than 1. To see this, note that the highest-MPC agents in the economy are the former employees of the shut down sector. They do not benefit from any government spending. They do benefit from direct transfers, but none of their spending will return to them as income. Thus, the typical Keynesian-cross amplification is broken as the highest-MPC agents in the economy do not benefit from spending by households or the government that was induced by fiscal policy.

We next extend the model to consider the effect of business closings. To do so, we

---

[1]Rowe (2020) provides a colorful and careful intuition for this possibility.

introduce a continuum of varieties agents can consume. Each variety is produced by a business that is required to pay a fixed cost to remain open and operational. In this case, we identify a *firm exit multiplier.* As some initial set of businesses is shut down (e.g. restaurants), e.g. due to health concerns, Keynesian forces reduce demand for other businesses (e.g. car dealers) as well. This, however, might mean that some of the other businesses become unable or unwilling to remain open. When they close, however, laying off their workers, a new, endogenous, Keynesian supply shock is born, that amplifies the existing exogenous one, creating a multiplier effect that may be sufficiently strong to shut down most of the economy.

We use this framework with endogenous business operations to discuss the effectiveness of several policies. We find that profit subsidies or employer-side payroll tax cuts are effective in keeping businesses afloat and preventing closures. Crucially, however, we find that these policies only work because they are conditional on remaining open; a lump-sum transfer to businesses would not necessarily prevent closures. When fixed costs stem from debt obligations, we find that monetary policy adopts a new transmission channel in the recession. By lowering debt payments, it can help prevent businesses from closing.

A related model features labor hoarding in the sector hit by a productivity shock. Firms may engage in labor hoarding, holding on to their workers at a loss, or let them go at the given wage, destroying the match and losing on future profits. When the incentives to keep workers is not large enough the analysis replicates our earlier results. However, for low enough interest rates firms put enough weight on future profits relative to current losses and decide to keep and pay their workers. In our model, this results in perfect insurance, possibly solving the demand deficiency at given interest rates. The decision to destroy job-worker matches may have longer run consequences on productivity, making the recovery after the shock more difficult. If firms are liquidity constrained then this distorts their labor hoarding decisions, so policies that mitigate these liquidity problems or improve firm balance sheets, while providing an incentive to keep workers, may improve the outcome.

Finally, we study the jointly optimal health and macroeconomic policy caused by a pandemic. We nest our previous model in a setting where we can model the health concerns more explicitly, private and social, and think about optimal policy, both of the Pigouvian nature and the macro stabilization. We show that the first best policy in our model involves closing down contact-intensive sectors and insurance payments to affected workers.

Literature Motivated by Pandemic

A large set of papers has emerged and is still expanding on macroeconomic issues surrounding the COVID-19 pandemic. There are a number of policy proposals, with a large number of them collected in Baldwin and Weder di Mauro (2020).

Fornaro and Wolf (2020) consider a standard New Keynesian representative-agent economy and study a pandemic as a negative shock to the growth rate in productivity. They also consider endogenous technological change and stagnation traps. In contrast, we focus on temporary shocks to supply due to shutdowns.

Faria e Castro (2020) builds on studies different forms of fiscal policy in a calibrated DSGE New Keynesian model. The model builds on Faria e Castro (2018) and features incomplete markets in the form of borrowers and savers with financial frictions. The pandemic is modeled as a large negative shock to the utility of consumption. Our paper instead focuses on a supply shock, motivated by the shutdowns, and studies the induced effects on demand.

A growing number of recent papers, motivated by the recent COVID-19 pandemic, make contact with epidemiological SIR or SIER models of contagion, merging them into an economic setting.[2] Atkeson (2020) provides a useful overview of the epidemiological models and their implications in the current COVID-19 pandemic. Berger, Herkenhoff and Mongey (2020) present an extended model with immunity and random testing. Eichenbaum, Rebelo and Trabandt (2020) consider a real one-sector dynamic model analysis and studies the effect of the pandemic taking into account optimal rational responses by private agents. They then consider optimal Pigouvian policy to internalize the externalities. Alvarez et al. (2020) study the optimal dynamic shutdown policy within a canonical SIR model. None of these papers focus on demand shortages or feature multiple sectors.

Jorda et al. (2020) provide some time-series evidence from historical pandemics on the impact on rates of return. The pandemics they study are persistent, with large numbers casualties.They find evidence that pandemics reduce the real rate of interest. It is not clear if this is comparable to the events we focus on, since we do not focus on the longer-term effects of death, but instead on the shorter-term effects of shutdowns that respond to the pandemic.

---

[2]Of course, a larger prior literature in history, health and development economics studied pandemics, and just to name a few recent examples, Philipson (1999), Greenwood et al. (2019) and Fogli and Veldkamp (2020).

2   Single Sector: Standard Supply Shocks Even with Incomplete Markets

We being by studying the effects of a supply shock in a one-sector model. We find that supply shocks have standard features here: they never cause demand effects strong enough to dominate the effects on the supply side. This applies even in economies with heterogenous agents and incomplete markets.

The framework we use for this section, and that we expand on in later sections, is a standard infinite horizon model with a single good. The model is populated by a unit mass of agents whose preferences are represented by the utility function

$$\sum_{t=0}^{\infty} \beta^t U(c_t), \tag{1}$$

where c_t is consumption and $U(c) = c^{1-\sigma}/(1-\sigma)$ is a standard CES utility function with intertemporal elasticity (EIS) σ^{-1}. Each agent is endowed with $\overline{n} > 0$ units of labor which are supplied inelastically. Competitive firms produce the final good from labor using the linear technology

$$Y_t = N_t.$$

The supply shock we introduce in the economy is inspired by the recent COVID-19 epidemic: a random fraction $\phi > 0$ of agents is unable to go to work in period $t = 0$. This captures the idea that the epidemic is making it unsafe for some agents to work, e.g. because their job requires close interaction with the public, so these agents stay home, either by choice or due to government containment policies.[3] Thus these agents can no longer supply their labor endowments in the first period. Starting in $t = 1$, we assume that all agents can again supply their full labor endowments of $\overline{n}$.

We analyze the effects of this supply shock separately for two versions of the model; first with complete markets, that is, with a representative agent; then with incomplete markets. In both cases, we look for two indicators—the response in the (natural) interest rate and the response of output if the real interest rate does not (or cannot) adjust in line with the natural rate. These indicators reveal whether the supply shock has standard effects or Keynesian effects. In the first case, the natural interest rate increases, and aggregate demand falls less than aggregate supply at a fixed real interest rate. In the second, the

---

[3]For now we take this as given. Section 6 studies an extension of the model with contagion, including both private and public (externalities) motives. We find that it can be optimal to shut down parts of the economy.

natural interest rate falls, and aggregate demand falls more than aggregate supply at a fixed real rate.

2.1   Complete Markets

Consider first the case of complete markets and thus a representative agent. Although the argument here is well known, it is useful to review it to set the stage for the rest of the analysis. Since markets are complete, we can view the supply shock as reducing the representative agent's labor supply from $\bar{n}$ to $(1 - \phi)\bar{n}$ in period $t = 0$.

What happens to the natural interest rate? Consider the flexible price version of this economy in which labor is always fully employed. The effect of the labor supply shock is mechanical: consumption falls at $t = 0$ before returning to its previous level. Thus, at date $t = 0$, the real interest rate rises to

$$1 + r_0 = \frac{1}{\beta} \frac{U'\left((1 - \phi)\bar{n}\right)}{U'\left(\bar{n}\right)} > \frac{1}{\beta}$$

above its previous steady state level of $1/\beta$.

The fact that the natural interest rate increases in this economy is a sign that there is no shortage of demand, in fact the opposite. To corroborate this logic, we introduce nominal rigidities. A convenient and tractable way to do so is to assume that nominal wages W_t are downwardly rigid. In that case, if labor demand falls below the labor endowment, wages are unchanged. This means that this economy can in principle display unemployment. We continue to assume that firms are perfectly competitive, so nominal prices are equal to nominal wages, $P_t = W_t$, and the real wage is $w_t = 1$.

In this economy, demand falls less than supply. To see why, let us do the following experiment. Assume that the central bank ensures full employment at all future dates so $c_t = \bar{n}$ for $t = 1, 2, ...$ Assume also that at $t = 0$ the central bank tries to keep the real interest rate at its steady state level $1/\beta - 1$. Consumption is then purely determined by the forward looking condition

$$U'\left(c_0\right) = \beta \frac{1}{\beta} U'\left(\bar{n}\right),$$

which yields $c_0 = \bar{n}$. This means that aggregate demand is completely unaffected, while aggregate supply falls to $(1 - \phi)\bar{n}$.[4] We summarize the results with complete markets.

---

[4]This cannot be an equilibrium, so the real rate will have to raise to its natural rate. For our purposes here, we do not need to specify the mechanism by which equilibrium is restored. A way of thinking about

Proposition 1. *Consider the single-sector model with complete markets. The negative supply shock causes an increase in the natural interest rate. If the real rate does not adjust, there is excess demand in the labor market.*

An alternative interpretation of Proposition 1 is that of a positive news shock. At $t = 0$, the representative agent learns that labor endowments will increase over time, from $(1 - \phi)\bar{n}$ to $\bar{n}$ per period. This intuitively explains why labor demand exceeds supply at $t = 0$ unless the real interest rate adjusts.

2.2 Incomplete Markets

Next we move to an economy with incomplete markets. The effects of the supply shock are less obvious here. After all, the agents hit by the shock lose their earnings and, in the presence of market incompleteness, might severely cut back their spending. We will start with a relatively simple setting with incomplete markets, but the results apply to more general setups.

For this economy, we label agents by $i \in [0, 1]$. Each agent i maximizes utility (1) subject to the budget constraint

$$c_{it} + a_{it} \leq w_t n_{it} + (1 + r_{t-1})a_{it-1}$$

Here, we assume agents have access to real, zero-net-supply, one-period bonds, paying real interest rate r_t. For the mass $1 - \phi$ of non-shocked agents, labor supply n_{it} equals $\bar{n}$. For the mass ϕ of shocked people, $n_{it} = 0$. We assume that market incompleteness takes the form of a borrowing constraint,

$$a_{it} \geq 0 \tag{2}$$

which is imposed on a fraction $\mu \in [0, 1]$ of agents. Henceforth, we refer to such agents as "constrained" agents. There is the same fraction μ of constrained agents in both sectors. The limit case $\mu = 0$ yields the same outcome as the representative agent complete-markets case; the limit case $\mu = 1$ corresponds to the case in which (2) is imposed for all agents. The economy starts in the symmetric steady state, in which $a_{i,-1} = 0$ for all agents i.

Once again, we begin by characterizing the flexible price equilibrium and deriving the response of the natural interest rate. Constrained agents hit by the shock see their income drop to zero and, due to (2), their consumption falls to zero as well, $c_{it} = 0$. All other agents

---

adjustment is to think that "off equilibrium" the excess demand of labor causes nominal wages to increase and the central bank responds to observed inflation by raising rates.

are on their Euler equation, that is,

$$U'(c_{i0}) = \beta(1 + r_0) U'(c_{i1}).$$

Due to homothetic preferences, their total consumption, which we denote by $\mathbf{c}_t$, also satisfies the Euler equation,

$$U'(\mathbf{c}_0) = \beta(1 + r_0) U'(\mathbf{c}_1). \tag{3}$$

Moreover, the goods market clearing condition has to hold in each period, implying that in all periods $t > 0$,

$$\mathbf{c}_t + \mu\phi\overline{n} = \overline{n} \tag{4}$$

and at date $t = 0$,

$$\mathbf{c}_0 = (1 - \phi)\overline{n}. \tag{5}$$

To understand condition (4), note that $\mu\phi\overline{n}$ is the steady-state spending of constrained agents that were shocked at date 0. As their consumption falls to zero at date 0, it is missing in (5). Substituting (4) and (5) into the Euler equation (3), we arrive at the expression for the natural interest rate,

$$1 + r_0^* = \frac{1}{\beta}\frac{U'((1 - \phi)\overline{n})}{U'((1 - \mu\phi)\overline{n})} \geq \frac{1}{\beta}. \tag{6}$$

Once more, the natural interest rate increases in response to the shock, as $(1 - \phi)\overline{n} \leq (1 - \mu\phi)\overline{n}$. The logic is similar to the one above. All agents on their Euler equation would like to smooth their consumption over time. Yet, that would lead to demand $\mathbf{c}_0$ that exceeds labor supply $(1 - \phi)\overline{n}$. Thus, the natural rate increases. What is new in the incomplete markets model is that not all agents are on their Euler equation and thus are able to smooth consumption. This reduces the necessity of interest rates to rise. In fact, in the special case where all agents are subject to the borrowing constraint (2), $\mu = 1$, the interest rate remains unchanged.

Now turn to the economy with rigid nominal wages, in which the central bank again attempts to implement a fixed interest rate $1 + r_0 = 1/\beta$. By (3) and (5), this implies that agents on their Euler equation demand

$$\mathbf{c}_0 = (1 - \mu\phi)\overline{n}$$

exceeding the supply of labor $(1 - \phi)\overline{n}$. In particular, each non-shocked agent would have to supply $n_0 = \frac{1 - \mu\phi}{1 - \phi}\overline{n} \geq \overline{n}$ units of labor. Again, the supply shock leads to a boom in labor

demand for each non-shocked worker, unless $\mu \to 1$.

The case of $\mu \to 1$ turns out to be an instructive extreme case. With $\mu \to 1$, each shocked agent that loses the income $\bar{n}$ cuts back spending by exactly $\bar{n}$, that is, the agents responds with a marginal propensity to consume (MPC) of 1. Thus, taken together, the shock removes $\phi\bar{n}$ units of labor supply and $\phi\bar{n}$ units of labor demand. On net, therefore, labor market clearing still holds without any other agent changing their behavior, and without interest rates moving.

This intuition turns out to greatly transcend the specific model written here. In fact, it applies to much richer incomplete-markets models, possibly with uninsurable idiosyncratic shocks, as long as they still have a single final good. The result being that in *none* of these models can the natural interest rate fall in response to the supply shock. To see this, note that the "best case scenario" for a falling interest rate is one in which all shocked agents cut back their spending 1-for-1 with the income shock they suffered. Yet, in that scenario, the supply shock reduces demand by exactly as much as it reduces supply. Thus, the interest rate does not move. Since this was the "best case scenario", it follows, that when shocked agents do not respond 1-for-1 to their income shocks, the natural rate always increases.

We summarize the insights from the incomplete markets model.

Proposition 2. *Consider the single-sector model with incomplete markets. The negative supply shock causes an increase in the natural interest rate. When real rates do not or cannot adjust, this translates into excess demand (a boom) in the labor market. In the corner case in which shocked agents cut their spending one for one with their income ($\mu \to 1$), the natural interest rate remains constant.*

3   Multiple Sectors: Keynesian Supply Shocks

We now enrich the model to include more than one sector. For some of the supply shock examples we would like our model to apply to, this is very natural. For example, the spread of the COVID-19 pandemic and the associated containment policies have clearly had asymmetric effects on different sectors. Particularly affected have been service sectors that require personal contact between consumers and workers.

This section will thus work with two sectors, 1 and 2. Later on, we will allow for a continuum of sectors. We assume that a fraction ϕ of agents works in sector 1, and a fraction $1 - \phi$ of agents works in sector 2. As before, agents inelastically supply labor $\bar{n}$ to their respective sector in the flexible price equilibrium; they may supply less than $\bar{n}$

in the equilibrium with wage rigidities. For now, we assume that workers are perfectly specialized in their sector.

The supply shock in this section will be one that prevents sector 1 agents from working. In terms of the COVID-19 example, consumption and production in sector 1 may require consumers and producers to meet in person. Consumption and production in sector 2, however, can take place without any personal contact. Containment measures may then be aimed at limiting contagion by preventing sector 1 agents from working.[5]

The technology to produce both goods is linear

$$Y_{jt} = N_{jt},$$ (7)

for $j = 1, 2$. Competitive firms in sector j hire workers at the sector-specific wage W_{jt} and sell good j at price P_{jt}. Prices P_{jt} are flexible and, given the technology above, the price of good j will be $P_{jt} = W_{jt}$.

Consumer preferences are now represented by the utility function

$$\sum_{t=0}^{\infty} \beta^t U\left(c_{1t}, c_{2t}\right),$$ (8)

where

$$U\left(c_{1t}, c_{2t}\right) = \frac{1}{1 - \sigma} \left(\phi^\rho c_{1t}^{1-\rho} + (1 - \phi)^\rho c_{2t}^{1-\rho}\right)^{\frac{1-\sigma}{1-\rho}},$$

so the utility function features constant elasticity $1/\rho$ between the two goods and constant intertemporal elasticity of substitution $1/\sigma$. To ensure that the model is well behaved under our supply shock, which prevents sector 1 agents from working, we assume for now that $\rho < 1$. We will discuss later how relaxing CES preferences may be useful in a context in which consumption of a set of goods goes to zero.

Next, we characterize the response of this multi-sector economy analogously to how we studied the single-sector model in the previous section, beginning with the complete markets case.

---

[5]We study the optimal containment policy in Section 6.

3.1 Complete Markets

Consider first the economy in steady state, before the shock hits, assuming all prices adjust flexibly so the economy reaches full employment. The equilibrium allocation is

$$c_1^* = Y_1^* = \phi \bar{n}, \quad c_2^* = Y_2^* = (1 - \phi)\bar{n}.$$

By symmetry, the relative price of good 1 in terms of good 2 is

$$p^* = 1. \tag{9}$$

The real interest rate is $1/\beta$ as in the one good economy, since consumption is constant in steady state. For reasons that will be clear shortly, it is useful to focus on the real interest rate in terms of good 2, defined as

$$1 + r_t \equiv (1 + i_t) \frac{P_{2t}}{P_{2t+1}}$$

where i_t denotes the nominal interest rate. The real interest rate $1 + r_t$ enters the Euler equation for good 2,

$$U_{c_2}(c_{1t}, c_{2t}) = \beta(1 + r_t)U_{c_2}(c_{1t+1}, c_{2t+1}), \tag{10}$$

where U_{c_j} denotes the partial derivative of U with respect to c_{jt}.

At date $t = 0$, when the supply shock hits, production in sector 1 shuts down, so

$$c_{10} = Y_{10} = n_{10} = 0.$$

Of course, there can no longer be full employment in sector 1. That is the inevitable effect of the shock. So we ask what happens in sector 2. As before, the shock is temporary and the economy goes back to steady state at $t = 1$. And, as before, we look first at what happens to the real interest rate to maintain full employment of the workers in sector 2; then we look at what happens to aggregate demand if the central bank keeps the real rate unchanged.

Using the representative agent's Euler equation (10), the natural rate after the shock is

$$1 + r_0 = \frac{1}{\beta} \frac{U_{c_2}(0, c_2^*)}{U_{c_2}(c_1^*, c_2^*)}. \tag{11}$$

The natural interest rate falls due to the epidemic shock if the ratio of marginal utilities on

13

Figure 2: When are supply shocks Keynesian with a representative agent?
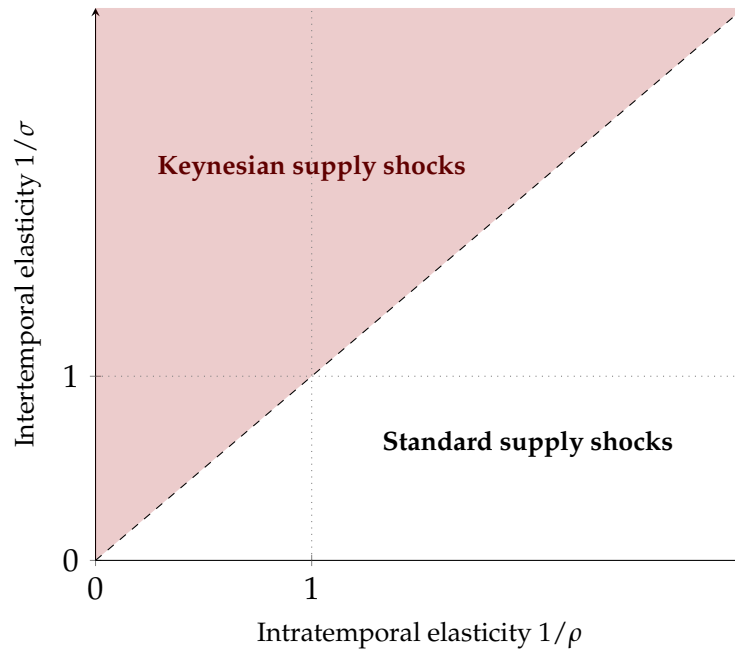
Note. We can extend the validity of (12) to the case $1/\rho < 1$ by replacing 0 in (11) with c_1 and letting $c_1 \to 0$.

the right-hand side is smaller than 1, or, using the functional forms introduced above, if

$$(1 - \phi)^{\frac{\rho - \sigma}{1 - \rho}} < 1.$$

An immediate consequence of this inequality is the following.

Proposition 3. *In the multi-sector model with complete markets, the negative supply shock translates into a reduced natural interest rate if and only if*

$$\frac{1}{\rho} < \frac{1}{\sigma}. \tag{12}$$

The interpretation of this result is straightforward. If the inequality (12) is satisfied the two goods are complements, so a drop in the production of good 1 increases the marginal utility of good 2, acting like a negative demand shock for good 2. To incentivize consumers to keep consuming enough of good 2 to keep employment at $\bar{n}$, we need a drop in the interest rate. We graphically illustrate the condition in Figure 2.

Before further discussing the plausibility of (12), we look at the effects of the shock on aggregate demand, assuming downwardly rigid nominal wages as before. A simple corollary of the proposition above is that if (12) is satisfied and the central bank keeps the

14

real interest rate at its steady state value $1/\beta - 1$, then there is an inefficient recession in sector 2, and the size of the output drop is given by

$$\frac{n_{20}}{\overline{n}} = (1 - \phi)^{\frac{1}{\sigma}\frac{\rho-\sigma}{1-\rho}}. \tag{13}$$

Thus, when $\rho > \sigma$ and the central bank does not (or cannot) act, the economy features two types of job losses: the unavoidable job losses $\phi\overline{n}$ due to the direct effect of the shock, and the inefficient job losses $\overline{n} - n_{20}$ due to insufficient demand in sector 2, with n_{20} given in (13).

Is $\rho > \sigma$ a plausible parameter configuration? If we look at standard models with a continuum of goods or varieties, the elasticity of substitution among goods $1/\rho$ is usually calibrated at values much larger than 1, while the intertemporal elasticity $1/\sigma$ is usually somewhere near 1. Such a choice of parameters would give us the opposite configuration. In that case, the two goods are substitutes, so a recession in sector 1 produces a demand boom in sector 2. Wages increase, generating inflation, and the central bank has to increase the nominal rate to avoid it.

One can argue whether that standard calibration is appropriate here, since here we are not interested in the elasticity of substitution among different varieties of the same good, but instead across different sectors, or across different locations. When restaurants are locked down in a whole city, it is hard to substitute them for restaurant meals elsewhere. This is why choosing an appropriate number for ρ requires thinking about whether the person-to-person services that get affected by containment policies are mostly complements or substitutes to other goods and services produced in the economy. In sum, we think condition (12) might be satisfied for the case of the recent COVID-19 supply shock.

What real interest rate? Before turning to incomplete markets, it is useful to provide an alternative interpretation of the result above. Notice that the shutdown of sector 1 can be interpreted as making the shadow price of good 1 prohibitively high. The ideal consumer price index (CPI) in this economy is

$$P_t = \left(\phi P_{1t}^{\frac{\rho-1}{\rho}} + (1 - \phi)P_{2t}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}.$$

If we set the price P_{1t} to infinity in period 0, the price index is still well defined, with $\rho < 1$. For a given nominal interest rate i_0, and assuming zero inflation in good 2, the real interest

rate in terms of the aggregate consumption basket C_t is

$$(1 + i_0) \frac{P_0}{P_1} = (1 + i_0)(1 - \phi)^{\frac{\rho}{\rho - 1}} > 1 + i_0.$$

Proceeding in this way and using the Euler equation $U'(C_0) = \beta(1 + i_0)\frac{P_0}{P_1}U'(C_1)$ it is then possible to re-derive equation (13). We can then reinterpret the shock causing unemployment in this economy as a shock that, for a given nominal rate, leads to a sharp temporary increase in the real interest rate, due to the fact that the shadow price of a number of goods goes up to infinity, as the goods cannot be bought. While this interpretation is useful, it is a bit harder to match to observables, because the shadow price of the goods not traded is not observed and their quantity goes to zero. So in the following we remain focused on what happens to the real interest rate in terms of the goods that are still traded. In other words, we focus on the *measured* CPI, rather than the ideal CPI.

The discussion above explains why we find it useful to work with the real interest rate in terms of the goods that are traded in all periods, which in this section means good 2.

3.2 Incomplete Markets

We now generalize this economy to allow for market incompleteness, exactly as in Section 2. Again, we use a simple description of market incompleteness but all results in this section generalize to richer setups, e.g. those in Werning (2015). In particular, a random fraction μ of households is subject to the borrowing constraint (2), and all households have the same initial financial wealth $a_{i0} = 0$.

To derive the response of the natural rate, we focus again on the group of agents who are not shocked or not constrained. Denote their consumption of goods 1 and 2 by $\mathbf{c}_{1t}$ and $\mathbf{c}_{2t}$, aggregated across the group. Due to homothetic preferences, we have Gorman aggregation. Thus, if their Euler equation holds individually, it also holds for the group, so

$$1 + r_0 = \frac{1}{\beta} \frac{U_{c_2}(0, \mathbf{c}_{20})}{U_{c_2}(\mathbf{c}_{11}, \mathbf{c}_{21})}. \tag{14}$$

To evaluate this expression, consider first the labor market clearing condition for sector 2 at date 0. Since shocked and constrained households consume nothing at date $t = 0$, labor market clearing requires

$$\mathbf{c}_{20} = (1 - \phi)\overline{n}.$$

At date $t = 1$, the group of not shocked or unconstrained household has total income

$(1 - \phi\mu)\, \bar{n}$ and consumes a fraction ϕ on good 1 and a fraction $1 - \phi$ on good 2.[6] We then have

$$\mathbf{c}_{11} = \phi\,(1 - \phi\mu)\,\bar{n}, \quad \mathbf{c}_{21} = (1 - \phi)\,(1 - \phi\mu)\,\bar{n}.$$

Substituting the consumption levels derived above, equation (14) yields

$$1 + r_0 = \frac{1}{\beta}(1 - \phi)^{\frac{\rho - \sigma}{1 - \rho}}\,(1 - \phi\mu)^\sigma.$$

Notice that the expression on the right-hand side is equal to $1/\beta$ when $\phi = 0$. Differentiating that expression with respect to ϕ and checking if the sign of the derivative is negative, yields the following result.

Proposition 4. *In the multi-sector model with incomplete markets, the negative supply shock translates into a reduced natural interest rate if and only if*

$$\frac{1}{\sigma} > \frac{1 - \mu}{1 - \phi\mu} \cdot \frac{1}{\rho} + \frac{\mu(1 - \phi)}{1 - \phi\mu}. \tag{15}$$

This result is similar to the one in Proposition 3, in that it provides a lower bound on the EIS for the natural rate to fall. In Proposition 3, that lower bound was given by the elasticity of substitution, $1/\rho$ (which is greater than 1). Proposition 4 shows that market incompleteness relaxes this condition, possibly considerably so. In particular, condition (15) only requires $1/\sigma$ to lie above a convex combination of $1/\rho$ and 1. Moreover, that convex combination converges to 1 as $\mu \to 1$. Thus, in the special case where the borrowing constraint applies to all agents, the condition for Keynesian supply shocks is simply given by

$$\frac{1}{\sigma} > 1,$$

or that the EIS exceeds 1. Interestingly, this condition no longer depends on the elasticity of substitution across goods, as long as $1/\rho > 1$. We illustrate condition (15) in Figure 3.

Before providing more of an intuition for this result, consider the case of a fixed real interest rate. If r_0 is fixed by the central bank at $1/\beta - 1$, the ratio of labor demand to labor

---

[6]The relative price of the two goods is always $p^* = 1$ in period 1, because this economy features Gorman aggregation and the wealth distribution does not affect relative prices. So individual consumption of all agents satisfy

$$\frac{c_{i11}}{c_{i21}} = \frac{\phi}{1 - \phi}\,(p^*)^{-\frac{1}{\rho}} = \frac{\phi}{1 - \phi}.$$

Figure 3: When are supply shocks Keynesian with incomplete markets?
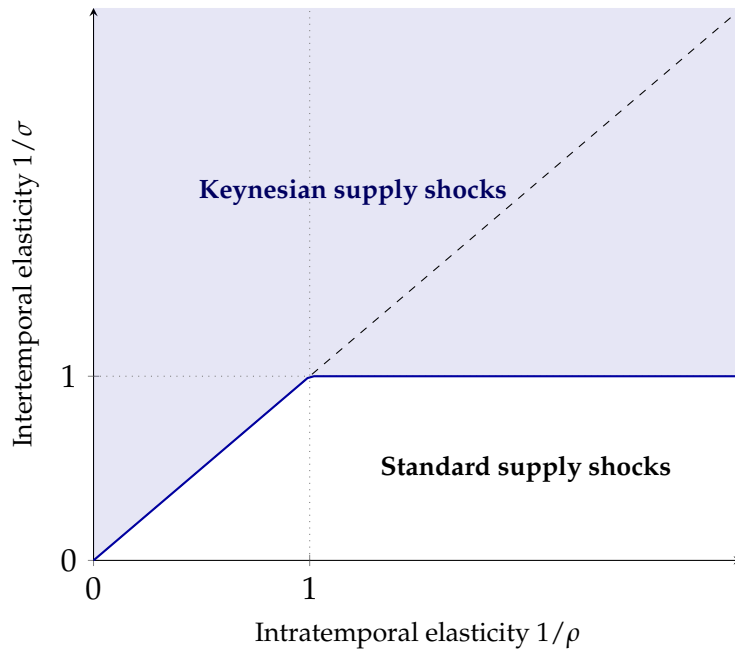
Note. We can extend (15) to the case $1/\rho < 1$ by replacing 0 in (14) with c_{10} and letting $c_{10} \to 0$. In that extension, we find that (15) becomes $1/\sigma \geq 1/\rho$ for $1/\rho < 1$.

supply in sector 2 is

$$\frac{n_{20}}{\overline{n}} = (1 - \phi\mu)(1 - \phi)^{\frac{1}{\sigma}\frac{\rho-\sigma}{1-\rho}}. \tag{16}$$

Under condition (15), the supply shock has Keynesian effects: labor demand falls below labor supply, causing a recession in the second sector.

Why does market incompleteness make it more likely for aggregate demand to fall? Compared to an economy with complete markets, a fraction μ of sector 1 agents cut their spending one-for-one with their income loss. This cut in spending weighs on aggregate demand above and beyond the spending response of unconstrained agents. Thus, aggregate demand falls more with incomplete markets.

We illustrate this in Figure 4 for a case where $1/\rho > 1/\sigma$. In that case, with complete markets, $\mu = 0$, sector 2 experiences a boom. With sufficient market incompleteness μ, however, condition (15) can be met, causing a bust in sector 2. Interestingly, in that case, a larger shock ϕ means a greater boom for small μ, but a greater bust for large μ.

No Paradox of Toil. An interesting perspective on this model is that it resolves New-Keynesian paradoxes at the ZLB. A number of papers have noticed that negative supply

Figure 4: Incomplete markets can cause recessions even with $1/\rho > 1/\sigma$.
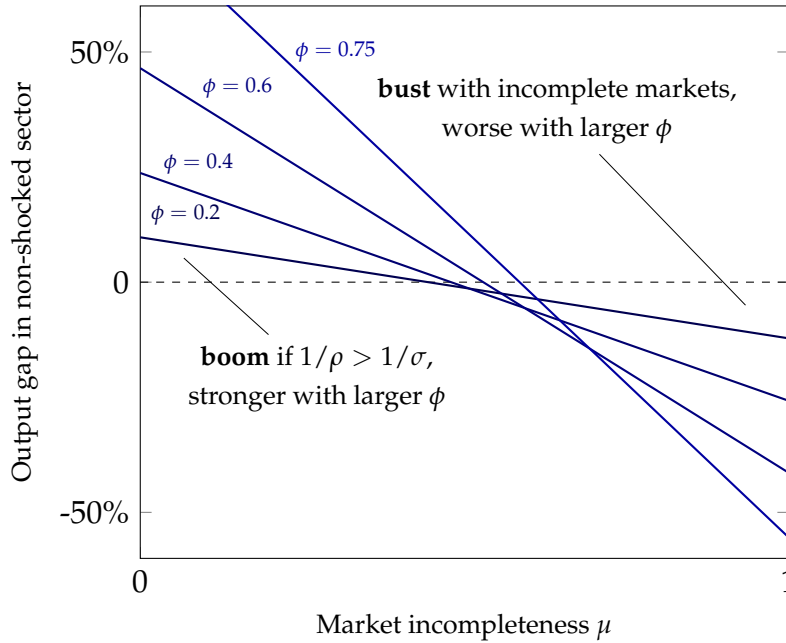
shocks, such as negative labor supply or negative TFP shocks, are expansionary at the ZLB in the New-Keynesian model (e.g. Eggertsson Paradox of Toil). This is not the case in our model if condition (15) is satisfied: In that case, output indeed falls in response to negative supply shocks, even when interest rates are fixed at the ZLB.

3.3 Fiscal Policy in the Incomplete Markets Model

One remedy against the ongoing COVID-19 recession that is currently being debated is fiscal stimulus. To consider the effects of fiscal stimulus in our model, we introduce a stylized government sector. We assume that the government chooses paths of government spending G_t, lump-sum transfers (or taxes) T_{jt} that can be targeted by sector, and government debt B_t, subject to the flow budget constraint

$$G_t + T_{1t} + T_{2t} + (1 + r_{t-1})B_t = B_{t+1}$$

We assume that in the steady state, $G = T_1 = T_2 = B = 0$.

We consider two stimulus policies. The first is traditional government spending, whereby the government raises G_0 at date zero, purchasing sector 2 goods, financed by uniform taxes in some future period $T_{1t} = T_{2t} < 0$. The second is a transfer program, such as unemployment insurance benefits, according to which the government chooses a

positive transfer $T_{1,0}$ to sector 1 consumers, again financed by uniform taxes in some future period.

Proposition 5. *Under government spending G_0 and transfers $T_{1,0}$, equilibrium employment in the incomplete markets is given by*

$$\frac{n_{20}}{\overline{n}} = \frac{G_0}{\overline{n}} + \mu \frac{T_{1,0}}{\overline{n}} + (1 - \phi\mu)(1 - \phi)^{\frac{1}{\sigma}\frac{\rho-\sigma}{1-\rho}}.$$

In particular, there is a unit government spending multiplier and a transfer multiplier equal to the average MPC. Both are smaller than predicted by the Keynesian cross.

This is a striking result. The average MPC in the economy is μ. In a typical recession in this model that affects both agents similarly (e.g. a discount factor shock) the multiplier would be $1/(1-\mu)$ and the transfer multiplier would be equal to $\mu/(1-\mu)$, exactly in line with the Keynesian cross (Galí et al. (2007), Farhi and Werning (2016), Auclert et al. (2018), Bilbiie (2019)). This is not the case here, however. The spending multiplier is simply 1, and the transfer multiplier is μ. Both multipliers are therefore missing the amplification through the Keynesian cross.

Why? The reason is that sector 1 is shut down: No agent can spend on sector 1. This means that any money spent by agents or the government flows into the pockets of sector 2 workers, and *not* sector 1 workers, who are up against their borrowing constraint and thus have greater MPCs. This suggests that traditional fiscal stimulus is less effective in a recession caused by our supply shock.

3.4   Labor Mobility

For now, we have assumed workers cannot move between sectors. It is easy to extend the analysis to the case in which they can move. In particular, suppose a fraction α of workers in each sector can move to the other sector. Consider first what happens with complete markets. Now the full employment level of output and consumption in sector 2 is larger as that sector can absorb the labor supply $\alpha\phi\overline{n}$ of the mobile workers initially in sector 1. The natural rate is now

$$1 + r_0 = \frac{1}{\beta} \frac{U_{c_2}(0, (1-\phi+\alpha\phi)\overline{n})}{U_{c_2}(c_1^*, c_2^*)} = \frac{1}{\beta}(1-\phi)^{\frac{\rho-\sigma}{1-\rho}} \left(\frac{1-\phi+\alpha\phi}{1-\phi}\right)^{-\sigma}.$$

The condition for the natural rate to fall is now weaker than with immobile labor. In particular, the natural rate now also falls in the case $\rho = \sigma$. The reason for this is simple:

sector 2 can now temporarily absorb some of the workers in sector 1, so spending in sector 2 needs to be temporarily larger. The consumption profile of good 2 is thus decreasing over time, requiring a lower rate.

If we look at employment under a constant real rate, the level of employment is still given by (13) (under $\rho > \sigma$) and is independent of α. This result may seem surprising, but simply follows from the purely forward looking nature of consumption decisions in the complete markets economy. Notice however that if we now look at the total employment losses in this economy

$$\bar{n} - (1 - \phi) n_{20} = \bar{n} - (1 - \phi) (1 - \phi)^{\frac{1}{\sigma} \frac{\rho - \sigma}{1 - \rho}} \bar{n},$$

their decomposition between efficient employment losses and excess employment losses changes. Efficient losses are now given by $\bar{n} - (1 - \phi + \alpha\phi) \bar{n}$, excess employment losses are $(1 - \phi + \alpha\phi) \bar{n} - n_{20}$. More mobility means that there are more excess losses.

Let us turn now to the incomplete market case. Now labor mobility affects directly spending decisions, for a given r_0, because the workers who can move do not lose their income. Now the fraction of workers who lose their income and are constrained becomes $(1 - \alpha) \mu\phi$, so if the interest rate is fixed, the employment losses in sector 2 are now

$$\frac{n_{20}}{\bar{n}} = (1 - (1 - \alpha) \phi\mu) (1 - \phi)^{\frac{1}{\sigma} \frac{\rho - \sigma}{1 - \rho}}.$$

Less mobility causes a deeper recession by causing larger income losses.

Excess employment losses in the incomplete market economy are

$$\left[(1 - \phi + \alpha\phi) - (1 - (1 - \alpha) \phi\mu) (1 - \phi)^{\frac{1}{\sigma} \frac{\rho - \sigma}{1 - \rho}} \right] \bar{n}.$$

If the following inequality holds

$$\mu(1 - \phi)^{\frac{1}{\sigma} \frac{\rho - \sigma}{1 - \rho}} < 1$$

it is still true that excess employment losses are larger in the economy with more mobility, as in the complete market case, because the effects on incomes that affects the demand side is weaker than the effect on the full employment output level.

3.5 Demand Chains

The analysis so far has focused on the degree of substitutability between the two goods for the possibility of Keynesian supply shocks. Introducing input-output relations allows us to look at the complementarity between the two sectors in a different light. The basic idea here is that restaurants (in sector 1) may need services from accountants (in sector 2) to produce their final good. When restaurants shut down that reduces a source of demand for accounting services. This logic suggests that input-output relations could increase the degree of complementarity between sectors, beyond the degree of complementarity driven purely by preferences.

Let us investigate this idea formally by introducing a simple input-output structure. In particular, we consider the possibility that goods produced by sector 2, that do not require personal contact, are used as intermediate inputs by sector 1, which requires personal contact.

The structure of preferences and markets is the same. We only change the technology. Good 1 is produced according to the production function

$$Y_1 = X^\alpha N_1^{1-\alpha},$$

where X is good 2 used as intermediate input in sector 1. The technology to produce good 2 is still linear

$$Y_2 = N_2.$$

Workers are still fully specialized with ϕ of them supplying $\bar{n}$ units of labor to sector 1 and $1 - \phi$ of them supplying $\bar{n}$ units of labor to sector 2. We use the term "demand chains" to capture the mechanism investigated here, because the usual argument is that supply chain disruptions cause an amplification of supply shocks upstream in the chain, while here we focus on disruptions happening downstream which reduce demand for upstream sectors.

First, consider the steady state economy before the shock. In steady state, the economy is at full employment and by market clearing

$$c_1^* = Y_1^* = X^{*\alpha}(\phi\bar{n})^{1-\alpha}, \quad c_2^* = Y_2^* = (1-\phi)\bar{n} - X^*,$$

where X^* is the steady state optimal level of intermediate input in sector 1. To find X^*, we just need the following two equations. The optimal demand for the intermediate input:

$$p\alpha X^{\alpha-1}(\phi\bar{n})^{1-\alpha} = 1,$$

22

where p is the relative price of good 1 in terms of good 2. And the relative demand for the two consumption goods, using market clearing at full employment:

$$p = \left(\frac{\phi}{1-\phi}\right)^\rho \left(\frac{X^\alpha(\phi\bar{n})^{1-\alpha}}{(1-\phi)\bar{n}-X}\right)^{-\rho}.$$

Substituting p from the second equation in the first, gives one equation in X, which can be shown to have a unique solution. The steady state real interest rate is $1/\beta - 1$ as before.

Consider now the effects of a temporary shutdown of sector 1. As before, this naturally implies that there cannot be full employment in the economy anymore. The difference is that now the shut down of sector 1 has an additional direct effect on the demand for good 2 due to the fact that there is no more demand for intermediate inputs for sector 1. As in the previous sections, let us first analyze what happens to the real interest rate if we want to keep full employment of the workers in sector 2. As in the model with no intermediate inputs, the natural rate after the shock that maintain full employment in sector 2 can be derived using the Euler equation in terms of the good 2

$$1 + r_0 = \frac{1}{\beta}\frac{U_{c_2}\left(0,(1-\phi)\bar{n}\right)}{U_{c_2}\left(c_1^*,c_2^*\right)}.$$

The natural rate falls after the shock if the ratio of marginal utilities is smaller than 1, which, using the functional forms we introduced is true if

$$\left(\frac{1-\phi}{p^*\left(x^*\right)^\alpha \phi^{1-\alpha}+1-\phi-x^*}\right)^{\frac{\rho-\sigma}{1-\rho}}\left(\frac{1-\phi-x^*}{1-\phi}\right)^{\rho+\rho\frac{\rho-\sigma}{1-\rho}} < 1$$

where $x^* \equiv X^*/\bar{n}$. This condition is weaker than the condition derived in Proposition 3, as for example it is satisfied when $\rho = \sigma$.[7] The reason is simple, in normal times there is a fraction x^* of labor supply in sector 2 is absorbed by the production of intermediates for sector 1. When the shock hits, this demand vanishes and needs to be replaced by direct demand by the consumers. This requires an increase in consumption of good 2, tilting the

---

[7]To make this argument precise, the condition is weaker if we calibrate the model with no intermediate inputs and the model with intermediate inputs so that the fraction of sector 2 in GDP is the same in steady state. That fraction is

$$1-\phi$$

in the baseline model and is

$$\frac{1-\phi}{p^*\left(x^*\right)^\alpha\phi^{1-\alpha}+1-\phi-x^*}$$

in the model with inputs, so the two models need different values of ϕ.

real rate downwards.

It is possible to extend the rest of the analysis. Skipping directly to what happens to output in the incomplete market case, when the central bank keeps the real rate at $1/\beta - 1$, we obtain that the recession in sector 2 is now given by

$$\frac{n_{20}}{\bar{n}} = (1 - \mu\phi) \left(\frac{1 - \phi}{px^\alpha \phi^{1-\alpha} + 1 - \phi - x} \right)^{\frac{1}{\sigma} \frac{\rho - \sigma}{1 - \rho}} \left(\frac{1 - \phi - x}{1 - \phi} \right)^{\frac{\rho}{\sigma} + \frac{\rho}{\sigma} \frac{\rho - \sigma}{1 - \rho}}.$$

The presence of the input-output structure adds the last factor in this expression, thus magnifying the effect of incomplete markets, given by the first factor.

4   Business Exit Cascades

An important problem in an economy hit by an adverse supply shock of the sort we consider in this paper is that businesses that are exposed to the shock see their revenues fall and might not be able to stay afloat. If businesses exit, however, workers will lose their jobs and might cut back on spending, feeding back into the magnitude of the recession. This feedback loop, and potential policies to break it, are investigated in the current section.

To do so, we generalize the model studied in the previous section. In particular, we now allow for a continuum of sectors $j \in [0, 1]$, each of which produces according to the linear aggregate production function (7). A representative worker supplies labor to each of the sectors, without cross-sectoral mobility. Preferences are still given by (8), but we change the consumption aggregator to

$$C_t = \left(\int c_{jt}^{1-\rho} dj \right)^{1/(1-\rho)}.$$

We assume that markets are incomplete in the sense used above. Specifically, we focus on the case where $\mu \to 1$, that is, all agents are subject to the borrowing constraint (2).

We assume that each sector j is monopolistically competitive, charging markups

$$\varphi \equiv \frac{1}{1 - \rho}.$$

Throughout this section, wages are assumed to be rigid, but prices are flexible. This implies that the real wage is constant at $w = 1 - \rho$ and profits of sector j are given by

$$\Pi_{jt} = \rho N_{jt}.$$

We make the simplifying assumption that projects earned by sector j are rebated lump-sum to the representative worker employed in sector j.[8]

We continue to study the same supply shock as before. In particular, a random fraction ϕ of agents can no longer supply labor to their sectors. Without loss, we assume these to be the sectors $j \in [0, \phi]$. Observe that, without any other changes, this model generates the exact same predictions as the model in Section 3.2, as we could bundle the sectors $j \in [0, \phi]$ together as "sector 1" and the remaining sectors as "sector 2". In particular, with ϕ inactive sectors, the date-0 response of employment n_0 in any of the active sectors is given by (16), that is,

$$\frac{n_0}{\overline{n}} = (1 - \phi)^{\frac{\rho}{\sigma} \frac{1-\sigma}{1-\rho}}. \tag{17}$$

4.1 The Business Exit Multiplier

We next allow for endogenous business exits. For this section, we define a business as owning a single sector j. Thus, businesses can also be labeled by j. Each period, a business has to pay a random fixed cost v_{jt}, which for simplicity we model as a transfer from business j to its representative worker. v_{jt} is drawn i.i.d. from a distribution $\Upsilon(v)$. We assume that $\Upsilon(0) = 0$ and $\Upsilon(\rho \overline{n}) = 1$. This ensures that no fixed cost realization leads to exit in the steady state.

Thus, business j makes profit $\Pi_{jt} - v_{jt}$ and finds it optimal to exit if $\Pi_{jt} < v_{jt}$. The mass of inactive businesses is denoted by $\hat{\phi}_t$. By definition, $\hat{\phi}_t$ always exceeds the fraction of shocked agents ϕ. Due to endogenous exit of businesses, however, $\hat{\phi}_t$ might be strictly greater than ϕ.

With this formulation, we have at date 0 that

$$1 - \hat{\phi}_0 = (1 - \phi) \, \Upsilon(\rho n_0). \tag{18}$$

Moreover, with $1 - \hat{\phi}_0$ active businesses, demand for employment is given by

$$\frac{n_0}{\overline{n}} = (1 - \hat{\phi}_0)^{\frac{\rho}{\sigma} \frac{1-\sigma}{1-\rho}} \tag{19}$$

Jointly, equations (18) and (19) pin down the mass of active businesses $\hat{\phi}_0$ as well as employment n_0.

The relationship between (18) and (19) is illustrated in Figure 5. The horizontal axis

---

[8]This assumption simplifies the algebra but does not materially affect the results.

Figure 5: The business exit multiplier

(a) Before the shock

(b) After the shock
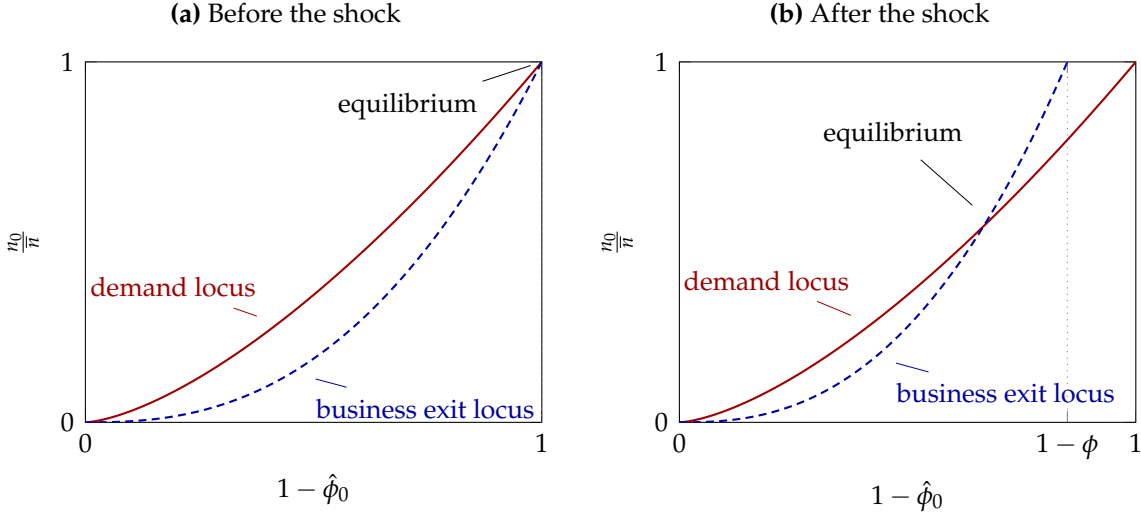
represents the mass of active businesses $1 - \hat{\phi}_0$. The vertical axis represents the employment relative to potential $n_0 / \bar{n}$. Under the assumption $\sigma^{-1} > 1$, both (18) and (19) describe positively sloped curves in Figure 5. We call (19) the "demand locus", as it describes the demand for employment, taking as given how many workers are able to work $(1 - \hat{\phi}_0)$. We call (18) the "(business) exit locus" as it describes the mass of businesses exiting given labor demand n_0.

When there is no shock, $\phi = 0$, the two curves necessarily intersect at coordinates $1 - \hat{\phi}_0 = 1$ and $n_0 / \bar{n} = 1$ (Panel a). However, a positive $\phi > 0$ shifts the exit locus to the left (Panel b). Interestingly, this shift raises the mass of inactive businesses by more than just ϕ, as additional workers laid off by exiting businesses also stop consuming. There is a cascade of business exits that generates a "business exit multiplier".[9]

A tractable special case. To see this even more cleanly, consider the following functional form for $Y(v)$,

$$Y(v) = \left(\frac{v}{\rho \bar{n}} \right)^\eta. \tag{20}$$

Here, $\eta > 0$ captures businesses' sensitivity to shutting down when average profits fall. A smaller η implies that almost all businesses have low fixed cost draws and thus stay in irrespective of profits. Vice versa, a larger η implies that businesses are more likely to exit when profits fall.

---

[9]Figure 5 shows that one can easily get multiple equilibria in this setting, when both curves intersect multiple times. We plan to investigate this case in future research.

With the functional form (20), the two equations (18) and (19) can be solved explicitly, giving

$$\log \frac{n_0}{\overline{n}} = \underbrace{\frac{1}{1 - \eta \frac{\sigma^{-1}-1}{\rho^{-1}-1}}}_{\text{firm exit multiplier}} \frac{\sigma^{-1}-1}{\rho^{-1}-1} \log(1-\phi). \tag{21}$$

This equation makes the business exit multiplier explicit. When $\eta = 0$, we are back in the case of Section 3.2 with exogenously active sectors. When $\eta > 0$, however, businesses' exit choices are endogenous to demand; but because supply shocks are Keynesian when $\sigma^{-1} > 1$, exit feeds back into less demand.

4.2 Policies

We use the framework laid out here to discuss the effectiveness of two policies.

Profit subsidy / employer-side payroll tax cut. The first policy is a profit subsidy, which, as in Section 3.3, we assume is paid for by employed agents. The subsidy raises profits by $1 + \tau$ for some $\tau > 0$. In our model, this is equivalent to an employer-side cut in payroll taxes. Such a subsidy enters (18), modifying it to

$$1 - \hat{\phi}_0 = (1-\phi) \, Y((1+\tau)\rho n_0)$$

and thus shifting the exit locus to the right. This mitigates some of the consequences of the shock. In the tractable special case studied above, the employment response is given by

$$\log \frac{n_0}{\overline{n}} = \frac{1}{1 - \eta \frac{\sigma^{-1}-1}{\rho^{-1}-1}} \frac{\sigma^{-1}-1}{\rho^{-1}-1} \left(\log(1-\phi) + \eta \log(1+\tau)\right).$$

Monetary policy. We model monetary policy as a change in the real interest rate $1 + r_0$ away from $1/\beta$. This clearly affects the demand locus (19), through the Euler equation. In particular, we have that

$$\frac{n_0}{\overline{n}} = (1 - \hat{\phi}_0)^{\frac{\rho}{\sigma} \frac{1-\sigma}{1-\rho}} \cdot (\beta(1+r_0))^{-1/\sigma}.$$

Accommodative monetary policy shifts the demand locus up, thereby also reducing the number of business exits in the economy.

Aside from the standard intertemporal substitution channel, however, there is another

transmission mechanism that can be active here. To illustrate this, we assume that when exiting, businesses sacrifice a claim to future profits Π, which instead is earned by new entrants in $t = 1$. This implies that the exit decision now compares current profits relative to fixed cost net of discounted future profits, $v_{jt} - \frac{1}{1+r_0}\Pi$. In this case, a "business exit channel" of monetary policy emerges. The exit locus now becomes

$$1 - \hat{\phi}_0 = (1 - \phi)\, Y \left(\rho n_0 + \frac{1}{1 + r_0}\Pi \right).$$

This channel operates by shifting the exit locus to the right.

5   Labor Hoarding vs Job Match Destruction

The previous analysis applied to businesses in the sectors that were not hit by the shutdown. We now briefly discuss how to model another margin, especially relevant for businesses in the sectors hit by the shock. To do so, we return to the two sector model from Sections 2-3, but make the following modifications.

5.1   A Simple Model of Labor Hoarding

Rather than describe the model in detail, we sketch out the model ingredients and main ideas for now. Production in the sector hit by the shock is carried out by firm-worker match pairs. Each firm is matched with a single worker, from some previous search process, which we shall not presently model for simplicity. We assume that these workers have a previously established wage w. There are no fixed costs, only the wage bill. These firms are owned in equal proportion by all agents in the economy, in both sectors.

For now we assume that if these workers are let go then they are able to return to work in the next period at $t = 1$, but not at the same firm, they instead match costlessly with new firms created at $t = 1$. This is a simplifying assumption to ensure that output at $t = 1$ remains anchored. We discuss relaxing this assumption below.

Will a firm in the shutdown sector wish to maintain the match or let the worker go? A firm considers the present value of its profits at $t = 0$ to be

$$V_0 = \max\{-w + \frac{1}{R}V_1, 0\}$$

where V_1 is the given present value of profits from $t = 1$ onwards if they do not break up the match. We assume that V_1 is strictly positive, i.e. they expect $py - w > 0$ in future

period.[10]

5.2 Monetary Policy Implications

Now suppose parameters are such that at the initial interest rate $R = 1/\beta$ firms wish to let go of their workers, $-w + \frac{1}{R}V_1 < 0$. Then these firms will not be worth anything. The analysis from Section 3 then applies, creating a demand deficient recession.

Suppose instead R is lowered sufficiently, so that $-w + \frac{1}{R}V_1 = 0$, or still lower. Then firms will want to keep their workers. As a result, this outcome achieves perfect insurance across workers in the two sectors: they both have the same income and financial assets. Note that the lower interest rate required for $-w + \frac{1}{R}V_1 = 0$ may be above or below the natural rate of interest that ensures full employment in the unaffected sector. If the interest rate is below this natural rate, then the monetary easing required to maintain job matches goes beyond that for full employment in the unaffected sector. Otherwise, there is a divine coincidence of sorts and the first-best outcome is achieved.

5.3 Scarring Job Match Losses

What happens if workers that are let go at $t = 0$ cannot immediately find a new job at $t = 1$? To be concrete and consider a simple case: suppose no matches can be created at $t = 1$, but they can be costlessly created at $t = 2$. Then if the interest rate is low enough to make firing workers an equilibrium the economy will suffer a recession over both periods $t = 0, 1$, effectively prolonging the duration of the supply shock. In period $t = 0$ the supply shock is exogenous, but in period $t = 1$ it results from the loss of job matches. Through an expectations channel, this may also make the recession at $t = 0$ deeper. More generally, a vast empirical literature has documented the scarring effects of job losses.[11]

The assumption that matches cannot be created at $t = 1$ but can be costlessly recreated at $t = 2$ is extreme, but we expect similar conclusions in a more elaborate model of search and vacancies, where job matches are created in a costly and incremental manner over time.

---

[10]For simplicity we assume for now that V_1 is identical across firms, but one can make V_1 or w vary across firms to get a smoother response to shocks.

[11]To be sure, in the context of shutdowns, we do not know if the effect of job losses is as damaging as during regular economic downturns or massive layoffs at firms during normal times. It is possible that job matches can be partly re-established at the end of a shutdown. Most likely, reality is a mix, where layoffs contribute to some scarring but potentially less than the ones we can expect during regular times.

5.4 Liquidity Problems and Policy Proposals

What happens if firms are liquidity constrained? If firms have some finite amount of liquidity at their disposal, say, because they cannot borrow nor issue equity and have limited past accumulated profits at their disposal, then they no longer maximize the present value of profits in an unconstrained fashion. This distorts firm decisions towards laying workers off, since the current period loss cannot be financed.

In this case, policies that directly affect the liquidity of firms or that insure firms for their loss in revenue, may restore the preferable outcome. In the model, this could be accomplished by a transfer to firms. In practice, these policies could be implemented in a number of ways and through a combination of fiscal and monetary branches of the government.

This discussion lends support to policy proposals at at the outset of the economic crises in March 2020 generated by the COVID-19 pandemic in the US and Europe. For example, Hamilton and Veuger (2020) propose emergency loans for small and medium sized firms most affected by liquidity problems facilitated by the Fed, complimented with tax credits on the fiscal side. Saez and Zucman (2020) propose an ambitious insurance policy, or "buyer of last resort", whereby the government makes up for any loss in revenue by in effect buying up the missing demand. Even some policy proposals aimed at paying workers directly, through unemployment benefits, emphasize the importance of preserving matches. For example, Dube (2020) calls for incentivizing temporary layoffs, so called furloughs, and the use of worker-sharing provisions, to keeps workers on payroll and allows workers to return easily after the shutdowns.[12]

6 Optimal Combined Shutdown and Macro Policy

Up to now we have taken the supply shock as given, just assuming certain sectors were inactive because of a lockdown. We now nest our model in a setting where we model more explicitly the health concerns, both private and social, and think about optimal policy, both in terms of Pigouvian interventions and of macro stabilization.

To this end, let us modify the consumers' objective function to include a health component. To keep things simple, we assume the health component is additive and does not directly affect the consumers' capacity to work. In particular, we modify the two sector

---

[12]Giupponi and Landais (2018) provide some evidence on related policies in Europe and study optimal policy in a model of labor hoarding and work-sharing.

model of Section 3, introducing the utility function:

$$\sum_{t=0}^{\infty} \beta^t \left(U \left(c_{1t}, c_{2t} \right) + h_t \right),$$

where

$$h_t = H \left(c_{1t}, n_{1t}, Y_{1t}, \xi_t \right)$$

is the consumer's health. The parameter ξ_t is the underlying shock and can take two values: $\underline{\xi}$ in normal times and $\overline{\xi}$ when there is an ongoing epidemic. When $\xi_t = \underline{\xi}$, the function H is just a constant. When $\xi_t = \overline{\xi}$, the function H is decreasing in c_{1t}, n_{1t} and Y_{1t}. The idea is that agents have a higher probability of being infected if they consume more in sector 1, if they produce more in sector 1, and if aggregate activity is higher in sector 1. The variables c_{1t} and n_{1t} are chosen by individual consumers, while the level of activity in sector 1, Y_{1t}, is taken as given by individual consumers. The presence of Y_{1t} captures the basic externality of an epidemic: more interactions in sector 1 cause a faster spread of the epidemic and so increase the probability of being infected for each person. The rest of the model is identical to the model in Section 3.

As in the rest of the paper, we assume that the shock is temporary and unexpected, so $\xi_0 = \overline{\xi}$ and $\xi_t = \underline{\xi}$ for $t = 1, 2, \ldots$. We assume that a government lockdown makes it impossible to consume and produce good 1. Therefore, under a lockdown the equilibrium analysis Section 3 applies unchanged here, as h_t is constant, because either $\xi_t = \underline{\xi}$ or $\xi_t = \overline{\xi}$ and $h_0 = H \left(0, 0, 0, \overline{\xi} \right)$.

To discuss interactions between public health policies and macroeconomic policies, we begin by considering partial interventions and their effects, and then we work out a case in which both sets of policies are set optimally and achieve the first best allocation. Our results are organized around three remarks. We start with an elementary observation.

Remark 1. Involuntary unemployment is not necessarily socially inefficient in our model.

To make this point, consider what happens in an economy in which there is no containment policy in place, so both sectors are potentially active, despite the shock $\xi_0 = \overline{\xi}$. Even absent containment policies, private motives will still induce a contraction in activity in sector 1, as people try to avoid contagion by reducing consumption and labor supply. This contraction in activity may result in involuntary unemployment in sector 1. To show that in a simple case, consider the complete market economy with nominal wage rigidities. Suppose $\rho = \sigma$ and suppose the central bank keeps the interest rate unchanged, so sector 2

31

is at full employment and $Y_{2t} = (1 - \phi) \bar{n}$. We can have an equilibrium with

$$c_{10} = Y_{10} < \phi \bar{n},$$

if the following two conditions are satisfied

$$U_{c_1} \left(Y_{10}, (1 - \phi) \bar{n} \right) + H_{c_1} \left(Y_{10}, Y_{10}, Y_{10}, \bar{\bar{\xi}} \right) = U_{c_1} \left(c_1^*, c_2^* \right),$$

and

$$U_{c_1} \left(Y_{10}, (1 - \phi) \bar{n} \right) + H_{c_1} \left(Y_{10}, Y_{10}, Y_{10}, \bar{\bar{\xi}} \right) + H_{n_1} \left(Y_{10}, Y_{10}, Y_{10}, \bar{\bar{\xi}} \right) > 0. \tag{22}$$

The first condition is the Euler equation in terms of good 1 and it gives $Y_{10} < \phi \bar{n}$ simply because consumers try to avoid consuming good 1 due to $H_{c_1} < 0$. The second condition is the optimality condition for labor supply and implies that it is optimal for the consumers to supply $n_{10} = \bar{n}$ as the private benefit from consumption, captured by the first two terms, exceeds the private cost of working, captured by the last term. The expression (22) can be interpreted as a Keynesian wedge, as the only disutility from work in our model comes from health costs.

Once we take into account the public health aspect, the presence of unemployment may not be socially inefficient, as agents do not internalize the externality in H. That is, it is possible that

$$U_{c_1} \left(Y_{10}, (1 - \phi) \bar{n} \right) + H_{c_1} \left(Y_{10}, Y_{10}, Y_{10}, \bar{\bar{\xi}} \right) + H_{n_1} \left(Y_{10}, Y_{10}, Y_{10}, \bar{\bar{\xi}} \right) + H_{Y_1} \left(Y_{10}, Y_{10}, Y_{10}, \bar{\bar{\xi}} \right) < 0,$$

so reducing further activity in sector 1 increases social welfare. The last equation shows that the Keynesian wedge, captured by the first three terms, can be more than compensated by a Pigouvian wedge, captured by the last term.

In the example above, there is a trade-off between public health objectives and aggregate demand stabilization. That happens in an example in which there are no public health policies are in place. Once public health policies are introduced, in the form of a lockdown, are the social welfare benefits of macro stabilization larger? That is, are the two policies complementary? The next remark shows that in our context the answer is yes.

Remark 2. There are complementarities between public health policies and aggregate demand stabilization.

The basic reason for this remark is that public health policies can produce a Keynesian supply shock and macro policies can then be helpful to correct the effects of the latter.

Again consider the example above and now suppose the government shuts down sector

1. Consider again the complete market economy, with nominal rigidities. Suppose now that $\rho > \sigma$, so we have an inefficient recession in sector 2. Lowering r_0 (assuming we do not hit the ZLB) allows the government to reach a first best efficient allocation if the following condition is satisfied at the full employment allocation

$$U_{c_1}\left(0, (1-\phi)\,\bar{n}\right) + H_{c_1}\left(0,0,0,\overline{\overline{\xi}}\right) + H_{n_1}\left(0,0,0,\overline{\overline{\xi}}\right) + H_{Y_1}\left(0,0,0,\overline{\overline{\xi}}\right) < 0. \qquad (23)$$

This condition means that a corner solution with a complete shutdown of sector 1 is socially efficient, as the public health benefits are large enough, given the shock $\overline{\overline{\xi}}$.

What happens when markets are incomplete? Now the social planner has to take into account three possible sources of inefficiency: inefficiency due to the public health externality, inefficiency due to lack of insurance, inefficiency due to involuntary unemployment. The next remark shows that if the government has sufficient tools it can deal with all of them and restore first best efficiency. In discussing the remark, we will show that again there are relevant complementarities between the tools used. In particular, social insurance policies that intervene on the second inefficiency, can ameliorate the dilemma between the other two that can arise with incomplete markets.

Remark 3. In the incomplete markets economy, a combination of public health policies, social insurance policies, and monetary policy can achieve the first best for a utilitarian social planner.

For this example, we need the incomplete market version of the two sector model, with nominal wage rigidities. Suppose that parameters are such that a shutdown policy produces a Keynesian supply shock, so there is inefficient unemployment in sector 2. Suppose also that monetary policy is constrained by a ZLB constraint and suppose that this constraint is binding in equilibrium of the incomplete market economy, with a shutdown. And suppose also that the ZLB constraint is not binding with complete markets. We know such a configuration is possible possible by Proposition 4.

Suppose first that the only policy tools available are a lockdown and monetary policy, and monetary policy is stuck at the ZLB. Suppose we can relax slightly the containment policy and increase output in sector 1 by dY_1. Consider the marginal benefit of this increase, for a utilitarian social planner. The effects on the consumption component of utility is

$$\int_0^1 \left[U_{c_1}\left(0, c_{i10}\right) \partial c_{i10} + U_{c_2}\left(0, c_{i20}\right) \partial c_{i20}\right] di$$

where ∂c_{ij0} is the effect of dY_1 on the consumption of consumer i of good j, in general

equilibrium. This effect can be large when the fraction of constrained agents μ is large for two reasons: for distributional reasons, as there are consumers with zero consumption of both goods, and for the presence of inefficient involuntary unemployment in sector 2. So it is possible that relaxing the containment policy may be desirable, as a second best way of correcting these two inefficiencies.

Suppose now that the government can also introduce a social insurance policy that reallocates income from sector 2 workers to sector 1 workers, so as to equalize their after-transfer incomes. Given this policy, the constraint on monetary policy is no longer binding, as we are now effectively in the complete markets economy. Moreover, if condition (23) is satisfied, a complete shut down of Sector 1 is now optimal.

In the example just discussed, there is a combination of policies that achieves the first best allocation: a containment policy that shuts down sector 1, a social insurance policy that compensates the workers in sector 1, and a monetary policy that hits the natural rate. The fact that the social insurance policy makes it easier to achieve the demand stabilization objective is not surprising per se: it is an example of a fiscal policy that makes it easier to do monetary policy. The novel observation is that this type of fiscal policy also makes it less costly for the government to impose a larger supply shock on the economy, that is, it makes it easier to pursue public health objectives.

7   Concluding Remarks

This paper asks a simple question: can a shock to supply, such as those experienced during a pandemic, lead to deficient demand? What are the combination of policy tools, monetary and fiscal, that best address this question in our model? Our answer is positive, demand may indeed overreact to the supply shock and lead to a demand-deficient recession. We have tried to lay out the conditions for this to be the case. Low substitutability across sectors and incomplete markets, with liquidity constrained consumers, all contribute towards the possibility of Keynesian suppy shocks. We then showed that various forms of fiscal policy, per dollar spent, may be less effective in our model. Despite this, the optimal policy to face a pandemic in our model combines as loosening of monetary policy as well as abundant social insurance.

References

Alvarez, Fernando, David Argente, and Francesco Lippi, "A Simple Planning Problem

for COVID-19 Lockdown," March 2020. mimeo.

Atkeson, Andrew, "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios," Working Paper 26867, National Bureau of Economic Research March 2020.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub, "The intertemporal keynesian cross," Technical Report, National Bureau of Economic Research 2018.

Baldwin, Richard and Beatrice Weder di Mauro, *Economics in the Time of COVID-19,*, VoxEU CEPR Press, 2020.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "An SEIR Infectious Disease Model with Testing and Conditional Quarantine," Working Paper 2020-25, Becker Friedman Institute for Economics 2020.

Bilbiie, Florin O, "The new Keynesian cross," *Journal of Monetary Economics*, 2019.

Dube, Arin, "Filling the Holes in Family and Business Budgets: Unemployment Benefits and Work Sharing in the Time of Pandemics," *Economics for Inclusive Prosperity*, March 2020, (24).

Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt, "The Macroeconomics of Epidemics," Working Paper 26882, National Bureau of Economic Research March 2020.

Farhi, Emmanuel and Iván Werning, "Fiscal multipliers: Liquidity traps and currency unions," in "Handbook of Macroeconomics," Vol. 2, Elsevier, 2016, pp. 2417–2492.

Faria e Castro, Miguel, "Fiscal Multipliers and Financial Crises," Working Papers 2018-23, Federal Reserve Bank of St. Louis October 2018.

— , "Fiscal Policy during a Pandemic," Working Papers 2020-006, Federal Reserve Bank of St. Louis March 2020.

Fogli, Alessandra and Laura Veldkamp, "Germs, Social Networks, and Growth," *Review of Economic Studies (forthcoming)*, 2020.

Fornaro, Luca and Martin Wolf, "Covid-19 Coronavirus and Macroeconomic Policy," Technical Report, CEPR March 2020.

Galí, Jordi, J David López-Salido, and Javier Vallés, "Understanding the effects of government spending on consumption," *Journal of the european economic association*, 2007, *5* (1), 227–270.

Giupponi, Giulia and Camille Landais, "Subsidizing Labor Hoarding in Recessions: The Employment & Welfare Effects of Short Time Work," CEP Discussion Papers dp1585, Centre for Economic Performance, LSE December 2018.

Gourinchas, Pierre-Olivier, "Flattening the Pandemic and Recession Curves," 2020. mimeo Berkeley University.

Greenwood, Jeremy, Philipp Kircher, Cezar Santos, and Michele Tertilt, "An Equilibrium Model of the African HIV/AIDS Epidemic," *Econometrica*, 2019, *87* (4), 1081–1113.

Hamilton, Steven and Stan Veuger, "How to Help American Businesses Endure and Jobs Survive," 2020. mimeo George Washington University.

Jorda, Oscar, Sanjay S. Singh, and Alan M. Taylor, "Longer-run economic consequences of pandemics," March 2020. mimeo.

Philipson, Tomas, "Economic Epidemiology and Infectious Diseases," Working Paper 7037, National Bureau of Economic Research March 1999.

Rowe, Nick, "Blog Post: Relative supply shocks, Unobtainium, Walras' Law, and the Coronavirus," *Worthwhile Canadian Initiative A mainly Canadian economics blog*, March 2020.

Saez, Emmanuel and Gabriel Zucman, "Keeping Business Alive: The Government as Buyer of Last Resort," 2020. mimeo Berkeley.